

Dynamic Control of Tunable Sub-optimal Algorithms for Scheduling of Time-varying Wireless Networks

Mahdi Lotfinezhad, Ben Liang, Elvino S. Sousa

Abstract—It is well known that for ergodic channel processes the Generalized Max-Weight Matching (GMWM) scheduling policy stabilizes the network for any supportable arrival rate vector within the network capacity region. This policy, however, often requires the solution of an NP-hard optimization problem. This has motivated many researchers to develop sub-optimal algorithms that approximate the GMWM policy in selecting schedule vectors. One implicit assumption commonly shared in this context is that during the algorithm runtime, the channel states remain effectively unchanged. This assumption may not hold as the time needed to select near-optimal schedule vectors usually increases quickly with the network size. In this paper, we incorporate channel variations and the time-efficiency of sub-optimal algorithms into the scheduler design, to dynamically tune the algorithm runtime considering the tradeoff between algorithm efficiency and its robustness to changing channel states. Specifically, we propose a Dynamic Control Policy (DCP) that operates on top of a given sub-optimal algorithm, and dynamically but in a large time-scale adjusts the time given to the algorithm according to queue backlog and channel correlations. This policy does not require knowledge of the structure of the given sub-optimal algorithm, and with low overhead can be implemented in a distributed manner. Using a novel Lyapunov analysis, we characterize the throughput stability region induced by DCP and show that our characterization can be tight. We also show that the throughput stability region of DCP is at least as large as that of any other static policy. Finally, we provide two case studies to gain further intuition into the performance of DCP.

Index Terms—Throughput stability region, dynamic tuning, channel variation, approximate GMWM time-efficiency

I. INTRODUCTION

The problem of scheduling of wireless networks has been extensively investigated in the literature. A milestone in this context is the seminal work by Tassiulas and Ephremides [2], where the authors characterized the *network-layer capacity region* of constrained queueing systems, including wireless networks, and designed a *throughput-optimal* scheduling policy, commonly referred to as the GMWM scheduling. In this context, capacity region by definition is the largest region that can be stably supported using any policy, including those with the knowledge of future arrivals and channel states. A throughout-optimal policy is a policy that stabilizes the network for any input rate that is within the capacity region and, thus, has the largest stable throughput region. In general [3][4], the GMWM scheduling should maximize the sum of backlog-rate products at each timeslot given channel states, which

A short version of this submission is published in [1].

The authors are with the Department of Electrical and Computer Engineering, University of Toronto. E-mail: {mlotfinezhad, liang}@comm.utoronto.ca, es.sousa@utoronto.ca.

can be considered as a GMWM problem. This problem has been shown to be, in general, complex and NP-hard [5][4][6]. Even in those cases where the optimization problem can be solved polynomially, distributed implementation becomes a major obstacle. These issues, naturally, motivated researchers to study and develop suboptimal centralized or distributed algorithms that can stabilize a fraction of the network-layer capacity region [7][5][6][8][9].

One implicit but major assumption in this context is that the time required to find an appropriate scheduling vector, *search-time*, is negligible compared to the length of a timeslot, or otherwise, during this search-time, channel states remain effectively unchanged. Since many algorithms take polynomial time with the number of users to output a solution [5][6][9], we see that this assumption may not hold in practice for networks with large number of users. In particular, it is possible that once an optimal solution corresponding to a particular channel state is found, due to channel variations, it becomes outdated to the point of being intolerably far away from optimality.

Intuitively, for many suboptimal algorithms, the solution found becomes a better and more *efficient* estimate of the optimal solution as the number of iterations increases or more *time* is given to the algorithm, e.g., see PTAS in [6]. This inspires us to consider this *time-efficiency* correspondence as a classifying tool for sub-optimal algorithms. As mentioned earlier, however, the solution found might become outdated due to channel variations. This poses a challenging problem as how the search-time given to sub-optimal algorithms should be adjusted to ensure an efficient scheduling with a large stable throughput region when channels states are time-varying.

Our work in this paper addresses the above challenge by joint consideration of channel correlation and time-efficiency of sub-optimal algorithms. In particular, we propose a dynamic control policy (DCP) that operates on top of a given sub-optimal algorithm A , where the algorithm is assumed to provide an approximate solution to the GMWM problem. Our proposed policy dynamically tunes the length of scheduling frames as the search-time given to the algorithm A so as to maximize the time average of backlog-rate product, improving the throughput stability region. This policy does not require the knowledge of input rates or the structure of the algorithm A , works with a general class of sub-optimal algorithms, and with low-overhead can be implemented in a distributed manner. We analyze the performance of DCP in terms of its associated throughput stability region, and prove that this policy enables the network to support all input rates that are within θ_∞ -scaled version of the capacity region. The scaling factor θ_∞ is a

function of the interference model, algorithm A , and channel correlation, and we prove that in general this factor can be tight. We also show that the throughput stability region of DCP is at least as large as the one for any other static scheme that uses a fixed frame-length, or search-time, for scheduling.

As far as we are aware, our study is the first that jointly incorporates the time-efficiency of sub-optimal algorithms and channel variations into the scheduler design and stability region analysis. One distinguishing feature of our work, apart from its practical implications, is the use of a Lyapunov drift analysis that is based on a *random* number of steps. Therefore, to establish stability results, we use a method recently developed for Markov chains [10], and modify it such that it is also applicable to our network model.

The rest of this paper is organized as follows. We review the related work in the next section. Network model including details of arrival and channel processes is presented in Section III. Structures of the sub-optimal algorithms and DCP policy are discussed in Section IV. We then provide performance analysis and the related discussion in Section V, followed by two case studies in Section VI. Finally, we conclude the paper in Section VII.

II. RELATED WORK

Previous work on throughput-optimal scheduling includes the studies in [2][11][3][12]. In particular, in [2], Tassiulas and Ephremides characterized the throughput capacity region for multi-hop wireless networks, and developed the GMWM scheduling as a throughput-optimal scheduling policy. This result has been further extended to general network models with ergodic channel and arrival processes [3]. Due to its applicability to general multi-hop networks, the GMWM scheduling has been employed, either directly or in a modified form, as a key component in different setups and many cross-layer designs. Examples include control of cooperative relay networks [12], rate control [13], energy efficiency [14][15], and congestion control [16][17]. This scheduling policy has also inspired pricing strategies maximizing social welfare [18], and fair resource allocation [16].

Another example of the throughput optimal control is the exponential rule proposed in [11]. In addition to the exponential rule scheduling, there are other approaches that use queue backlog, either explicitly or implicitly, for scheduling [19][20][21]. For instance, in [19], active queue management is used that implements CSMA protocol with backlog dependent transmission probabilities. It is shown that such an approach can implement a distributed fair buffer. In one other work [20], an adaptive CSMA algorithm is proposed that iteratively adjust nodes' aggressiveness based on nodes' (simulated) queue backlog.

The GMWM scheduling despite its optimality, in every timeslot, requires the solution of the GMWM problem, which can be, in general, NP-hard and Non-Approximable [6]. Thus, many studies have focused on developing sub-optimal constant factor approximations to the GMWM scheduling. One interesting study addressing the complexity issue is the work in [22], where sub-optimal algorithms are modeled as randomized algorithms, and it is shown that throughput-optimality can

be achieved with linear complexity. In a more recent work [23], the authors propose distributed schemes to implement a randomized policy similar to the one in [22] that can stabilize the entire capacity region. These results, however, assume non-time-varying channels. Other recent studies in [4][24] generalize the approach in [22] to time-varying networks, and prove its throughput-optimality. This optimality, as expected, comes at the price of requiring excessively large amount of other valuable resources in the network, which in this case is memory storage. Specifically, the memory requirement in [4][24] increases exponentially with the number of users, making the generalized approach hardly amenable to practical implementation in large networks.

Another example of sub-optimal approximation is the work in [5], where the authors assume that the controller can use only an *imperfect* scheduling component, and as an example they use maximal matching to design a distributed scheduling that is within a constant factor of optimality. This scheduling algorithm under the name of *Maximal Matching* (MM) scheduling and its variants have been widely studied in the literature [7][6][25][9][26][27]. In [7][5], it is shown that under simple interference models, MM scheduling can achieve a throughput (or stability region) that is at least half of the throughput achievable by a throughput-optimal algorithm (or the capacity region). Extended versions of these results for more general interference models are presented in [6][9], where in [9] randomized distributed algorithms are proposed for implementing MM scheduling, being a constant factor away from the optimality. This result has been further strengthened recently [28] stating that the worst-case efficiency ratio of Greedy Maximal Matching scheduling in geometric network graphs under the κ -hop interference model is between $1/6$ and $1/3$. All of the mentioned proposals so far either do not consider channel variations, or assume the search-time is relatively small compared to the length of a timeslot.

The closest work to ours in this paper is [8], where based on the linear-complexity algorithm in [22], the impact of channel memory on the stability region of a general class of sub-optimal algorithms is studied. Despite its consideration for channel variations, this work still does not model the search-time, and implicitly assumes it is negligible.

In this paper, we consider the problem of scheduling from a new perspective. We assume a sub-optimal algorithm A is given that can approximate the solution of the GMWM problem, and whose efficiency naturally improves as the search-time increases. We then devise a dynamic control policy which tunes the search-time, as the length of scheduling frames, according to queue backlog levels in the network, and also based on channel correlations. As far as we are aware, our study is the first that explicitly models the time-efficiency of sub-optimal approaches, and uses this concept along with channel correlation in the scheduler design.

III. NETWORK MODEL

We consider a wireless network with N one-hop source-destination pairs, where each pair represents a data flow¹.

¹ Extension to multi-hop flows is possible using the methods in [2][3].

Associated with each data flow, we consider a separate queue, maintained at the source of the flow, that holds packets to be transmitted over a wireless link. Examples of this type of network include downlink or uplink of a cellular or a mesh network.

A. Queueing

We assume the system is time-slotted, and channels hold their state during a timeslot but may change from one timeslot to another. Let $\mathbf{s}(t)$ be the matrix of all channels states from any given node i to any other node j in the network at time t . For instance, when the network is the downlink or uplink of a cellular network, $\mathbf{s}(t)$ will reduce to the vector of user-base-station channel states, i.e., $\mathbf{s}(t) = (s_1(t), \dots, s_N(t))$, where $s_i(t)$ is the state of the i th link (corresponding to the i th data flow) at time t . Throughout the chapter, we use bold face to denote vectors or matrices. Let \mathcal{S} represent the set of all possible channel state matrices with finite cardinality $|\mathcal{S}|$. Let $D_i(t)$ denote the rate over the i th link corresponding to the i th data flow at time t , and $\mathbf{D}(t)$ be the corresponding vector of rates, i.e., $\mathbf{D}(t) = (D_1(t), \dots, D_N(t))$. In addition, let $I_i(t)$ represent the amount of resource used by the i th link at time t , and $\mathbf{I}(t)$ be the corresponding vector, i.e., $\mathbf{I}(t) = (I_1(t), \dots, I_N(t))$. The vector $\mathbf{I}(t)$ contains both scheduling and resource usage information, and hereafter, we refer to it simply as the schedule vector. Let \mathcal{I} denote the set containing all possible schedule vectors, with finite cardinality $|\mathcal{I}|$.

Note that the exact specification of the scheduling vector $\mathbf{I}(t)$ is system dependent. For instance, in CDMA systems, it may represent the vector of power levels associated with wireless links; in OFDMA systems, it may represent the number of sub-channels allocated to each physical link; and when interference is modeled as the K-hop interference model [6], the vector can be a link activation vector representing a sub-graph in the network. Assuming that transmission rates are completely characterized given channel states, the schedule vector, and the interference model, we have

$$\mathbf{D}(t) = \mathbf{D}(\mathbf{s}(t), \mathbf{I}(t)).$$

We assume that transmission rates are bounded, i.e., for all $\mathbf{s} \in \mathcal{S}$ and $\mathbf{I} \in \mathcal{I}$,

$$D_i(\mathbf{s}, \mathbf{I}) < D_{max}, \quad 1 \leq i \leq N,$$

for some large $D_{max} > 0$.

Let $A_i(t)$ be the number of packets arriving in timeslot t associated with the i th link (or data flow), and $\mathbf{A}(t)$ be the vector of arrivals, i.e., $\mathbf{A}(t) = (A_1(t), \dots, A_N(t))$. We assume arrivals are i.i.d.² with mean vector

$$\mathbb{E}[\mathbf{A}(t)] = \mathbf{a} = (a_1, \dots, a_N),$$

and bounded above:

$$A_i(t) < A_{max}, \quad 1 \leq i \leq N,$$

for some large A_{max} .

²This assumption is made to simplify the analysis, and our results can be extended to non i.i.d arrivals.

Finally, let $\mathbf{X}(t) = (X_1(t), \dots, X_N(t))$ be the vector of queue lengths, where $X_i(t)$ is the queue length associated with the i th link (or data flow). Using the preceding definitions, we see that $\mathbf{X}(t)$ evolves according to the following equation

$$\mathbf{X}(t+1) = \mathbf{X}(t) + \mathbf{A}(t) - \mathbf{D}(t) + \mathbf{U}(t),$$

where $\mathbf{U}(t)$ represents the wasted service vector with non-negative elements; the service is wasted when in a queue the number of packets waiting for transmission is less than the number that can be transmitted, i.e., when $X_i(t) < D_i(t)$.

B. Channel State Process

We assume the channel state process is stationary and ergodic. In particular, for all $\mathbf{s} \in \mathcal{S}$, as $k \rightarrow \infty$, we have

$$\frac{1}{k} \sum_{i=0}^{k-1} \mathbf{1}_{\mathbf{s}(t+i)=\mathbf{s}} \rightarrow \pi(\mathbf{s}), \quad a.s.,$$

where $\mathbf{1}_{(\cdot)}$ denotes the indicator function associated with a given event, and $\pi(\mathbf{s})$ is the steady-state probability of state \mathbf{s} . Let \mathcal{P}_t represent the past history of the channel process and be defined by $\mathcal{P}_t = \{\mathbf{s}(i); 0 \leq i \leq t\}$. The above almost surely convergence implies that for any $\epsilon > 0$ and $\zeta > 0$, we can find a sufficiently large $K_{\epsilon, \zeta, t} > 0$ such that [29]

$$P\left(\sup_{k > K_{\epsilon, \zeta, t}} \left| \frac{1}{k} \sum_{i=0}^{k-1} \mathbf{1}_{\mathbf{s}(t+i)=\mathbf{s}} - \pi(\mathbf{s}) \right| > \epsilon \mid \mathcal{P}_t\right) < \zeta. \quad (1)$$

We assume that the almost surely convergence is *uniform* in the past history and t in the sense that regardless of \mathcal{P}_t and t , there exists a $K_{\epsilon, \zeta}$ such that (1) holds with $K_{\epsilon, \zeta, t} = K_{\epsilon, \zeta}$ ³.

C. Capacity Region

In our context, capacity region, denoted by Γ , is defined as the closure of the set of all input rates that can be *stably* supported by the network using any scheduling policy including those that use the knowledge of future arrivals and channel states. In [2][30] and recently under general conditions in [3], it has been shown that the capacity region Γ is given by

$$\Gamma = \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \text{Convex-Hull}\{\mathbf{D}(\mathbf{s}, \mathbf{I}) \mid \mathbf{I} \in \mathcal{I}\}.$$

IV. DYNAMIC CONTROL POLICY

As mentioned in the introduction, DCP controls and tunes the search-time given to a sub-optimal algorithm to improve the stability region. The considered sub-optimal algorithms are assumed to provide a sub-optimal solution to the GMWM problem. In the following, we first elaborate on the structure of the sub-optimal algorithms, and then, describe the operation of DCP.

³Examples of this channel model include but are not limited to Markov chains.

A. Sub-optimal Algorithms Approximating GMWM Problem

It is well known that the GMWM scheduling is throughput-optimal in that it stabilizes the network for all input rates interior to capacity region Γ . This policy in each timeslot uses the schedule vector $\mathbf{I}^*(t)$ that is argmax to the following GMWM problem:

$$\max \sum_{l=1}^N X_l(t) D_l(\mathbf{s}(t), \mathbf{I}), \quad \text{subject to } \mathbf{I} \in \mathcal{I}. \quad (2)$$

However, as mentioned in Section I, this optimization problem can be in general NP-hard. We therefore assume that there exists an algorithm A that can provide suboptimal solutions to the max-weight problem given in (2). To characterize the structure of algorithm A , let $\mathbf{I}^*(\mathbf{X}, \mathbf{s})$ be the argmax to (2) by setting $\mathbf{X}(t) = \mathbf{X}$ and $\mathbf{s}(t) = \mathbf{s}$. Thus,

$$\mathbf{I}^*(\mathbf{X}, \mathbf{s}) = \underset{\mathbf{I} \in \mathcal{I}}{\operatorname{argmax}} \mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}),$$

where $\mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I})$ is the scalar product of the two vectors, and for ease of notation, we have dropped the transpose symbol required for $\mathbf{D}(\mathbf{s}, \mathbf{I})$. In the rest of this paper, we use the same method to show the scalar products. Associated with $\mathbf{I}^*(\mathbf{X}, \mathbf{s})$, let $\mathbf{D}^*(\mathbf{X}, \mathbf{s})$ be defined as

$$\mathbf{D}^*(\mathbf{X}, \mathbf{s}) = \mathbf{D}(\mathbf{s}, \mathbf{I}^*(\mathbf{X}, \mathbf{s})). \quad (3)$$

Thus, $\mathbf{D}^*(\mathbf{X}, \mathbf{s})$ is the optimal rate, in the sense of (2), when the backlog vector is \mathbf{X} and the channel state is \mathbf{s} .

Let $\mathbf{I}^{(n)}$ be the *output* schedule vector of algorithm A when it is given an amount of time equal to n timeslots, $\mathbf{X}(t) = \mathbf{X}$, and $\mathbf{s}(t) = \mathbf{s}$. We therefore assume that the time given to algorithm A can be programmed or tuned as desired, or simply, the algorithm can continue or iterate towards finding better solutions over time. We assume that $\mathbf{I}^{(n)}$ is in general a random vector with distribution $\mu_{\mathbf{X}, \mathbf{s}}^{(n)}$. Since the objective function in (2) is a continuous function of $\mathbf{X}(t)$, we naturally assume that algorithm A characterized by the distribution of $\mathbf{I}^{(n)}$, for all $n \geq 1$, and all values of \mathbf{X} and \mathbf{s} , has the following property:

Assumption 1: For all $\mathbf{I} \in \mathcal{I}$, $\mathbf{s} \in \mathcal{S}$, and n , we have that

$$|\mu_{\mathbf{X}_1, \mathbf{s}}^{(n)}(\mathbf{I}^{(n)} = \mathbf{I}) - \mu_{\mathbf{X}_2, \mathbf{s}}^{(n)}(\mathbf{I}^{(n)} = \mathbf{I})| \rightarrow 0,$$

as $\mathbf{X}_1 \rightarrow \mathbf{X}_2$. In addition, assuming and keeping $\|\mathbf{X}_1 - \mathbf{X}_2\| < C$ for a given $C > 0$, the above convergence also holds when $\|\mathbf{X}_1\| \rightarrow \infty$. Moreover, the convergence becomes equality if $\mathbf{X}_1 = \beta \mathbf{X}_2$, for some $\beta > 0$.

In the following, we discuss concrete models that provide further details on the structure of algorithm A . Note that these models serve only as examples, and our results do not depend on any of these models; what required is only Assumption 1.

The first model arises from the intuition that the distribution $\mu_{\mathbf{X}, \mathbf{s}}^{(n)}$ should improve as n increases. More precisely, we can define the sequence $\{\mu_{\mathbf{X}, \mathbf{s}}^{(n)}, n = 1, 2, 3, \dots\}$ to be an *improving* sequence if for all $n > 1$,

$$\mathbb{E}[\mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}^{(n)})] \geq \mathbb{E}[\mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}^{(n-1)})] \geq \dots \geq \mathbb{E}[\mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}^{(1)})].$$

The first model uses the above and defines a *natural* algorithm to be the one for which the above inequalities hold for all values of \mathbf{X} and \mathbf{s} .

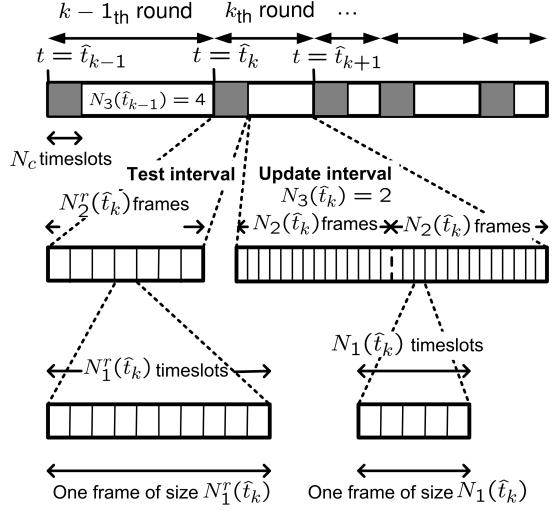


Fig. 1. Illustration of scheduling rounds, test intervals, update intervals, and frames.

As for the second model, we may have that $\mathbf{I}^{(n)}$ is such that

$$\mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}^{(n)}) \geq g(n) \mathbf{X} \mathbf{D}(\mathbf{s}, \mathbf{I}^*(\mathbf{X}, \mathbf{s})), \quad (4)$$

where the function $g(n)$ is a non-decreasing function of n , and less than or equal to one. For instance, if the optimization problem can be approximated to a convex problem [31], then $g(n) = \xi(1 - \zeta^n)$, where $0 < \xi \leq 1$ and $0 \leq \zeta < 1$. Another possible form for $g(n)$ is

$$\left(1 - \beta \frac{\ln N}{\ln n}\right),$$

where β is a positive constant. This form of $g(n)$ may stem from cases where the optimization problem associated with (2) admits Polynomial-Time Approximation Scheme (PTAS) [6].

The last model that we consider is a generalization of the previous model, where we assume that (4) holds with probability $h(n)$ as a non-decreasing function of n . This specification can model algorithms that use randomized methods to solve (2), and without its consideration for the improvement over n , is similar to the ones developed in [22][8].

B. Dynamic Control Policy and Scheduling

The dynamic control policy in this paper interacts with scheduling component, and through some measures, which will be defined later, dynamically tunes the time spent by the scheduler, or more precisely algorithm A , to find a schedule vector. In what follows, we describe the joint operation of DCP and the scheduler.

As DCP operates, the time axis becomes partitioned to a sequence of *scheduling rounds*, where each round might consist of a different number of timeslots. An illustrative example is provided in Fig. 1. Let \hat{t}_k denote the start time of the k th round. Each round begins with a *test interval* followed by an *update interval*. In the beginning of the test interval of each round, a *candidate* value for the number of timeslots given to the algorithm A to solve (2) is selected by DCP. Let $N_1^r(\hat{t}_k)$ denote this candidate value for the k th round, and

assume $N_1^r(\hat{t}_k) \in \mathcal{N}_1$, where \mathcal{N}_1 has a finite cardinality. In the rest, we use N_1^r instead of $N_1^r(\hat{t}_k)$ where appropriate. The algorithm that chooses the candidate value might be in general a randomized algorithm. Thus, we use the superscript r to make this point clear. We assume N_1^r takes an optimal value with probability $\delta > 0$, where optimality will be defined later by (7) and its following discussion.

We set the length of the test interval to be

$$N_1^r N_2^r = N_c = \text{const.},$$

a multiple of N_1^r , where N_2^r is adjusted accordingly so that the test interval has a fixed length N_c . Therefore, given N_1^r , the test interval becomes partitioned into N_2^r consecutive frames of N_1^r timeslots. In the beginning of each frame, e.g., at time t , the current backlog vector $\mathbf{X}(t)$ and channel state $\mathbf{s}(t)$ are provided to the algorithm A . The algorithm then spends N_1^r timeslots to find a schedule vector. Depending on the properties of a particular instance of algorithm A , this vector is used by the scheduler to update scheduling decisions in the next frame in a variety of methods.

In the first method, the schedule vector found after N_1^r timeslots in the frame starting at time t is used throughout the next frame of N_1^r timeslots starting at time $t + N_1^r$. Thus, the schedule vector used in any frame is obtained by using backlog and channel state information at the beginning of its previous frame. This method is general and can be applied to all types of algorithm A .

We can apply a second method where algorithm A is capable of outputting schedule vectors in intermediate steps, and not only after the planned N_1^r timeslots. Consider the i th timeslot of a given frame of N_1^r timeslots started at time t , where $i \leq N_1^r$. Suppose \mathbf{I}_{t+i}^c is the intermediate solution found by the algorithm A after i timeslots in the considered frame, and \mathbf{I}_p is the vector found at the end of its previous frame. Then, we may assume that with some probability, \mathbf{I}_{t+i}^c is used if

$$\mathbf{X}(t+i)\mathbf{D}(\mathbf{s}(t+i), \mathbf{I}_{t+i}^c) > \mathbf{X}(t+i)\mathbf{D}(\mathbf{s}(t+i), \mathbf{I}_p),$$

otherwise \mathbf{I}_p is used in the timeslot following the i th timeslot. The update rule in [8] provides an example where two schedule vectors are compared, and the best is selected with a well-defined probability.

As for the third method, we may assume algorithm A can accept an initial schedule vector. In this case, we can assume that the algorithm A at a given frame accepts the schedule vector found in the previous frame as the initial point to the optimization problem of (2). Note that many graph-inspired algorithms do not start from a given initial vector (as a sub-graph), but instead, gradually progress towards a particular solution. These algorithms⁴, therefore, do not belong to the class of algorithms considered for this method. A forth method can also be considered by mixing the second and the third method if algorithm A has the corresponding required properties. Our results in this paper extend to these methods as long as Property 1 and Property 2 in Section V-B hold.

⁴Adaption of these algorithm to time-varying networks is an interesting problem, and is left for the future research.

Given N_1^r , and a method to use the output of algorithm A , DCP evaluates scheduling performance resulting from the value for N_1^r . The performance criterion is the normalized time-average of the backlog-rate (scalar) product. To define the criterion precisely, let $\varphi(\cdot, \cdot, \cdot)$ be defined as

$$\varphi(t, n_1, n_2) = \sum_{j=0}^{n_2-1} \sum_{i=0}^{n_1-1} \frac{\mathbf{X}_{t+jn_1+i} \mathbf{D}_{t+jn_1+i}}{n_1 n_2 \|\mathbf{X}_t\|}.$$

If $\|\mathbf{X}_t\| = 0$, we set $\varphi(t, n_1, n_2) = 0$. Based on the above definition, the criterion associated with the test interval of the k th scheduling round, which is computed by DCP, is denoted by $\varphi^r(\hat{t}_k)$, where

$$\varphi^r(\hat{t}_k) = \varphi(\hat{t}_k, N_1^r(\hat{t}_k), N_2^r(\hat{t}_k)).$$

This quantity is then used to determine the length of frames in the update interval of the k th round.

Update intervals are similar to the test intervals in that they are consisted of a multiple number of fixed-length frames. More precisely, we assume that the update interval in the k th round becomes partitioned into $N_2(\hat{t}_k)N_3(\hat{t}_k)$ consecutive frames of $N_1(\hat{t}_k)$ timeslots. Integers $N_1(\hat{t}_k)$ and $N_2(\hat{t}_k)$ are such that

$$N_1(\hat{t}_k)N_2(\hat{t}_k) = N_c. \quad (5)$$

Therefore, the length of the k th update interval is $N_3(\hat{t}_k)$ times the length of a test interval. Moreover, we see that $N_1(\hat{t}_k)$ in the k th update interval takes the role of $N_1^r(\hat{t}_k)$ in the k th test interval. Assuming the same method is applied to all test and update intervals to use the output of algorithm A , we can properly define $\varphi(\hat{t}_k)$ as

$$\varphi(\hat{t}_k) = \varphi(\hat{t}_k + N_c, N_1(\hat{t}_k), N_2(\hat{t}_k)N_3(\hat{t}_k)).$$

The quantity $\varphi(\hat{t}_k)$ is similar to $\varphi^r(\hat{t}_k)$, and measures the normalized time-average of backlog-rate product in the k th update interval.

DCP, on top of algorithm A , uses $\varphi(\hat{t}_{k-1})$ and $\varphi^r(\hat{t}_k)$ to dynamically control the value of $N_1(\hat{t}_k)$ and $N_3(\hat{t}_k)$ over time. Specifically, in the k th round, at the *end* of the test interval, the policy chooses either the N_1 used in the previous update interval, $N_1(\hat{t}_{k-1})$, or the newly chosen value of N_1 in the current test interval, $N_1^r(\hat{t}_k)$, according to the following update rule:

$$N_1(\hat{t}_k) = \begin{cases} N_1^r(\hat{t}_k) & \text{if } \varphi^r(\hat{t}_k) > \varphi(\hat{t}_{k-1}) + \alpha \\ N_1(\hat{t}_{k-1}) & \text{otherwise} \end{cases},$$

where α is a suitably small but otherwise an *arbitrary* positive constant. At the same time, the value of $N_3(\hat{t}_k)$, is updated according to the following:

$$N_3(\hat{t}_k) = \begin{cases} \max(1, \frac{N_3(\hat{t}_{k-1})}{2}) & \text{if } \varphi^r(\hat{t}_k) > \varphi(\hat{t}_{k-1}) + \alpha \\ \min(L_1, 2N_3(\hat{t}_{k-1})) & \text{otherwise,} \end{cases}$$

where L_1 is a suitably large but otherwise an *arbitrary* positive constant. Note that $N_2(\hat{t}_k)$ becomes updated such that (5) holds. Once the values of N_1 , N_2 , and N_3 are updated, in the rest of the scheduling round, which by definition is the update interval, the policy proceeds with computing the time average $\varphi(\hat{t}_k)$. When the k th round finishes, the $k+1$ th round

starts with a test interval, and DCP proceeds with selecting $N^r(\hat{t}_{k+1})$, and applying the update rule at the end of the $k+1$ th test interval. This completes the description of joint operation of DCP and the scheduling component.

Considering the above description, we see that DCP keeps trying new values for N_1 . Once a good candidate is found for N_1 , the update rule with high probability uses this value for longer periods of time by doubling the length of update intervals. In case the performance in terms of the backlog-rate product degrades, the length of update intervals are halved to expedite trying new values for N_1 . Note that α can be arbitrarily small, but should be a positive number. This avoids fluctuations between different values of N_1 performing closely, thus preventing short update intervals. In addition, it limits incorrect favoring towards new values of N_1 in the test intervals, where due to atypical channel conditions, the normalized backlog-rate product deviates from and goes beyond its expected value. Finally, note that L_1 can be arbitrarily large, but should be a finite integer. This assumption is mainly analysis-inspired but is also motivated by the fact that a larger L_1 can lead to a larger delay.

V. PERFORMANCE ANALYSIS

In this section, we evaluate the performance of DCP in terms of its associated stability region. We first introduce several key definitions and functions, and then state the main theorem of the paper.

A. Definitions

Since the backlog vector is non-Markovian, we consider the following definition for the stability of a process.

1) *Stability*: Suppose there are a bounded closed region \mathcal{C} around the origin, and a real-valued function $F(\cdot) \geq 0$ such the following holds: For any t , and $\sigma_{\mathcal{C}}$ defined by

$$\sigma_{\mathcal{C}} = \inf\{i \geq 0 : \mathbf{X}_{t+i} \in \mathcal{C}\},$$

we have

$$\mathbb{E}[\sigma_{\mathcal{C}}] \leq F(\mathbf{X}(t)) \mathbf{1}_{\mathbf{X}(t) \notin \mathcal{C}}.$$

Then, the system is said to be stable.

This definition implies that when $\mathbf{X}(t) \notin \mathcal{C}$, e.g., when $\|\mathbf{X}(t)\|$ is larger than a threshold, the conditional expectation of the time required to return to \mathcal{C} , e.g., so that $\|\mathbf{X}(t)\|$ becomes less than or equal to the threshold, is bounded by a function of only $\mathbf{X}(t)$, uniformly in the past history and t . This definition further implies that if the sequence $\mathbf{X}(t)$ is stable, then [32]

$$\lim_{k \rightarrow \infty} \sup_t P(|\mathbf{X}(t)| > k) = 0.$$

2) *θ -scaled Region and Maximal Stability*: Suppose $0 \leq \theta \leq 1$. A region is called θ -scaled of the region Γ , and denoted by $\theta\Gamma$, if it contains all rates that are θ -scaled of the rates in Γ , i.e.,

$$\theta\Gamma = \{\mathbf{a}_1 : \mathbf{a}_1 = \theta\mathbf{a}_2, \text{ for some } \mathbf{a}_2 \in \Gamma\}.$$

Further, the θ -scaled region is called *maximally stable* if for all arrival rate vectors interior to $\theta\Gamma$, the system can be stabilized, and for all $\epsilon > 0$ there exists at least one rate vector interior to

$(\theta + \epsilon)\Gamma$ that makes the system unstable, both under the same given policy. Thus, maximal stability determines the largest *scaled* version of Γ that can be stably supported under a given policy.

B. Auxiliary Functions and Their Properties

To define the first function, hypothetically suppose for all t , $\mathbf{X}(t) = \mathbf{X}$ for a given \mathbf{X} , $\mathbf{X} \neq 0$, and thus, $\mathbf{X}(t)$ does not get updated. In addition, assume that N_1 has a fixed value over time. Considering these assumptions and an update interval of infinite number of frames⁵, each consisting of N_1 timeslots, we can see that in the *steady state*, the expected normalized backlog-rate product, averaged over one frame, is equal to

$$\phi(\mathbf{X}, N_1) = \mathbb{E}_{\mathbf{s}, A} \left[\frac{\sum_{i=1}^{N_1} \mathbf{X} \mathbf{D}_i}{N_1 \|\mathbf{X}\|} \right], \quad (6)$$

where \mathbf{D}_i is the rate vector in the i th timeslot of a given frame in the steady state. This expectation is over the steady-state distribution of channel process, and possibly over the randomness introduced by the algorithm A .

Intuitively, $\phi(\mathbf{X}, N_1)$ states how well a particular choice for N_1 performs, in terms of backlog-rate product, when queue-length changes are ignored. This is exactly what we need to study since the stability region often depends on the behavior of scheduling at large queue-lengths, where in a finite window of time the queue-lengths do not change significantly.

To simplify notation, where appropriate, we use t as the first argument of $\phi(\cdot, \cdot)$; by that we mean⁶

$$\phi(t, N_1) = \phi(\mathbf{X}(t), N_1).$$

Having defined $\phi(\mathbf{X}, N_1)$, we define $\tilde{N}_1(\mathbf{X})$ and $\tilde{\phi}(t)$ by

$$\tilde{N}_1(\mathbf{X}) = \operatorname{argmax}_{N_1 \in \mathcal{N}_1} \phi(\mathbf{X}, N_1), \quad (7)$$

and

$$\tilde{\phi}(t) = \tilde{\phi}(\mathbf{X}(t)) = \phi\left(\mathbf{X}(t), \tilde{N}_1(\mathbf{X}(t))\right).$$

Finally, for a given \mathbf{X} with $\|\mathbf{X}\| \neq 0$, we define

$$\chi(\mathbf{X}) = \mathbb{E}_{\mathbf{s}} \left[\frac{\mathbf{X} \mathbf{D}^*(\mathbf{X}, \mathbf{s})}{\|\mathbf{X}\|} \right],$$

where $\mathbf{D}^*(\mathbf{X}, \mathbf{s})$ is defined in (3), and the expectation is over the *steady-state* distribution of the channel process.

According to the above definitions, we see that when variations in the backlog vector are ignored after time t , and N_1 is confined to have a fixed value, $\tilde{N}_1(\mathbf{X}(t))$ becomes the optimal value for N_1 in terms of the normalized backlog-rate product, and $\tilde{\phi}(t)$ represents the corresponding expected value. In particular, note that $\tilde{N}_1(\mathbf{X})$ is a function of \mathbf{X} and may take different values for different \mathbf{X} 's. The quantity $\chi(\mathbf{X})$, on the other hand, is the expected normalized backlog-rate product if for all states we could find the optimal schedule vector. This quantity, therefore, can serve as a benchmark to measure performance of sub-optimal approaches.

⁵Here, we assume the channel evolves, and that the algorithm A is used in the same manner as it is used in an ordinary update interval with a finite N_c , as discussed in Section IV-B.

⁶By definition of $\phi(\cdot, \cdot)$, here we hypothetically assume the backlog vector $\mathbf{X}(t_1)$ for all times t_1 is equal to $\mathbf{X}(t)$.

Note that $\chi(\mathbf{X})$ is continuous function of \mathbf{X} and does not depend on $\|\mathbf{X}\|$. Similarly, by Assumption 1, $\phi(\mathbf{X}, N_1)$ does not depend on $\|\mathbf{X}\|$, and is expected to have the following property.

Property 1: Suppose $\|\mathbf{X}_1 - \mathbf{X}_2\| < C$ for a given $C > 0$. For any given $\epsilon > 0$, there exists a sufficiently large $M > 0$ such that if $\|\mathbf{X}_1\| > M$, then for all $N_1 \in \mathcal{N}_1$

$$|\phi(\mathbf{X}_1, N_1) - \phi(\mathbf{X}_2, N_1)| < \epsilon.$$

If the first or the second method in Section IV-B is used, this property holds since by Assumption 1, algorithm A statistically finds similar schedule vectors when two backlog vectors are close and large. In case the third or the forth method is used, it is possible to consider explicit restrictions for algorithm A such that $\phi(\mathbf{X}, N_1)$ is well-defined and Property 1 holds. However, in this paper, we simply assume that algorithm A is well-structured, in terms of the distribution of $\mathbf{I}^{(n)}$, so that by the ergodicity of the channel process this property also holds for these methods.

Recall that $\varphi^r(\hat{t}_k)$ is the normalized time-average of backlog-rate product over the k th test interval. If we assume that the backlog vector is kept fixed at $\mathbf{X}(\hat{t}_k)$, by ergodicity of the channel process as explained in Section III-B, we expect $\varphi^r(\hat{t}_k)$ to converge to $\phi(\hat{t}_k, N_1^r(\hat{t}_k))$. Hence, when the number of frames is large, which is the case when N_c is large, $\varphi^r(\hat{t}_k)$ should be close to $\phi(\hat{t}_k, N_1^r(\hat{t}_k))$ with high probability. However, the backlog vector is not fixed and changes over time. But by Assumption 1, algorithm A statistically responds similarly to different backlog vectors if they are close and sufficiently large. This can be exactly our case since arrivals and departures are limited, and thus, for a fixed N_c , the changes in the norm of backlog vector are bounded over one test interval. Therefore, by Assumption 1, if $\|\mathbf{X}(\hat{t}_k)\|$ is sufficiently large, the changes in the backlog have little impact on the distribution of $\varphi^r(\hat{t}_k)$. Applying a similar discussion to $\varphi(\hat{t}_k)$ while noting that the length of update intervals is bounded by $L_1 N_c$, we expect the following property.

Property 2: There exist $\varrho_\varphi > 0$ and $\theta_\varphi > 0$ such that for any given $\epsilon > 0$, there exists $M > 0$ such that if $\|\mathbf{X}_{\hat{t}_k}\| > M$, then regardless of k and the past history, up to and including time \hat{t}_k , with probability at least $(1 - \varrho_\varphi)$

$$|\varphi^r(\hat{t}_k) - \phi(\hat{t}_k, N_1^r(\hat{t}_k))| < \theta_\varphi + \epsilon.$$

Similarly, regardless of k and the past history, up to and including time $\hat{t}_k + N_c$, with probability at least $(1 - \varrho_\varphi)$

$$|\varphi(\hat{t}_k) - \phi(\hat{t}_k + N_c, N_1(\hat{t}_k))| < \theta_\varphi + \epsilon.$$

Moreover,

$$\lim_{N_c \rightarrow \infty} \varrho_\varphi = \lim_{N_c \rightarrow \infty} \theta_\varphi = 0.$$

According to the preceding discussion, we can see that θ_φ and ϱ_φ mainly measure how fast the time-averages converge to their expected value, and ϵ models the error due to variations in the backlog vector $\mathbf{X}_{\hat{t}_k+i}$. Thus, as stated above, ϱ_φ and θ_φ can be made arbitrarily small by assuming a sufficiently large value for N_c . In a practical implementation, however, N_c is a limited integer, and therefore, $\theta_\varphi > 0$ and $\varrho_\varphi > 0$. Note

that when the first or the second method in Section IV-B is used, Property 2 holds as a result of its preceding discussion, uniform convergence of the channel process, and finiteness of $|\mathcal{I}|$. Similar to Property 1, in the case of the third or the forth method, we assume this property results from the well-structuredness of algorithm A .

As the final step towards the main theorem, we define several random variables that are indirectly used in the theorem statement. Specifically, let i_δ be a geometric random variable with success probability δ' , where

$$\delta' = (1 - \varrho_\varphi)^2 \delta,$$

where δ is defined in Section IV-B. In addition, let i_φ be a r.v. with the following distribution.

$$P(i_\varphi = 0) = \varrho_\varphi,$$

and

$$P(i_\varphi = k) = (1 - \varrho_\varphi)^{2k-1} (1 - (1 - \varrho_\varphi)^2), \quad k \geq 1.$$

We also define the random sequence $\{N'_3(i), i \geq 1\}$ as⁷

$$N'_3(i) = \begin{cases} L_1 & (1 \leq i \leq i_\delta) \vee (i = i_\delta + i_\varphi + 1) \\ 1 & (i = i_\delta + 1) \wedge (i_\varphi = 1) \\ 2 & (i = i_\delta + 1) \wedge (i_\varphi > 1) \\ \min(\frac{2^i}{2^{i_\delta+2}}, L_1) & (i_\delta + 2 \leq i \leq i_\delta + i_\varphi) \wedge (i_\varphi > 1) \\ 0 & i > i_\delta + i_\varphi + 1 \end{cases}.$$

Using the above sequence, we define R_∞ as

$$R_\infty = \frac{\mathbb{E} \left[\sum_{i=i_\delta+1}^{i_\delta+i_\varphi} N'_3(i) \right]}{\mathbb{E} \left[\sum_{i=1}^{i_\delta+i_\varphi+1} (1 + N'_3(i)) \right]}, \quad (8)$$

which plays a key role in theorem statement and its proof. Note that for a fixed $\delta > 0$, we have

$$\lim_{\varrho_\varphi \rightarrow 0} R_\infty = \frac{L_1}{1 + L_1}.$$

As mentioned earlier, we can make ϱ_φ and θ_φ arbitrarily small by choosing a sufficiently large value for N_c . We are now ready to state the theorem.

C. Main Theorem on Stability of DCP

We have the following theorem:

Theorem 1: Consider a network as described in Section III. For this network, let θ be a constant defined by

$$\theta = R_\infty \inf_{\|\mathbf{X}\|=1} \frac{(\tilde{\phi}(\mathbf{X}) - \alpha - 3\theta_\varphi)}{\chi(\mathbf{X})}.$$

In addition, let θ_∞ be

$$\theta_\infty = \inf_{\|\mathbf{X}\|=1} \frac{\tilde{\phi}(\mathbf{X})}{\chi(\mathbf{X})}.$$

(a) If $6\theta_\varphi < \alpha$ and $2\alpha \leq \inf_{\|\mathbf{X}\|=1} \tilde{\phi}(\mathbf{X})$, then the network is stable under DCP if the mean arrival rate vector, \mathbf{a} , lies strictly inside the region $\theta\Gamma$.

⁷Here, \wedge and \vee are the *and* and *or* operators, respectively.

- (b) For any input rate strictly inside $\theta_\infty \Gamma$, there exist a sufficiently small value for α , and sufficiently large values for L_1 and N_c such that the network becomes stabilized under DCP. In other words, we can expand the sufficient stability region $\theta \Gamma$ arbitrarily close to $\theta_\infty \Gamma$ by choosing appropriate values for α , L_1 , and N_c .
- (c) There exist instances of networks, as described in Section III, for which their associated region $\theta_\infty \Gamma$ is maximally stable under DCP.

Proof: The proof is provided in the Appendix. ■

D. Discussion

1) *Intuitive Explanation of θ :* Theorem 1 states that all input rates interior to $\theta \Gamma$ can be stably supported under DCP. In particular, it implicitly quantifies θ as a function of the sub-optimality of algorithm A and channel state correlation. Clearly, the value of θ is not fixed, and can vary from a particular network setup to another. As expected, for a fixed \mathbf{X} , as algorithm A finds better schedule vectors in shorter times, and as the channel states become more correlated, $\tilde{\phi}(\mathbf{X})$ gets closer to $\chi(\mathbf{X})$, and θ gets closer to one, expanding the region $\theta \Gamma$ to the capacity region \mathcal{C} .

In addition, Theorem 1 shows how the stability region is directly affected by the choices for α and L_1 , and the values for θ_φ and ϱ_φ . The impact of α on θ could be predicted by noting that the update rule uses N_1^r in an update interval only when the normalized average backlog-rate product increases at least by α . Thus, we expect to see a decrease of the type $\frac{\alpha}{\chi(\mathbf{X})}$ in the stability region scaling. The effect of θ_φ and ϱ_φ is less obvious, but can be roughly explained as follows. Suppose at the k th round the optimal N_1 is selected, i.e., $N_1^r(\hat{t}_k) = \tilde{N}_1(\hat{t}_k)$. In this case, to have a proper comparison, $\varphi^r(\hat{t}_k)$ and $\varphi(\hat{t}_{k-1})$ should satisfy their corresponding inequalities in Property 2. Moreover, to make sure that $N_1^r(\hat{t}_k)$ or a near optimal N_1 is used in the l th round after the k th, we at least require $\varphi^r(\hat{t}_l)$ satisfy its corresponding inequality in Property 2. Therefore, there are at least three inequalities of the form in Property 2 that should be satisfied, which results in the term $3\theta_\varphi$ in the expression for θ .

The factor R_∞ in a sense measures the least fraction of time in update intervals where *near* optimal values for N_1 is used. To better understand R_∞ , suppose ϱ_φ is small, and the backlog vector is large. Once the optimal value for N_1 is found in a round, as long as the inequalities in Property 2 hold for the subsequent rounds, N_1 gets updated for only a few times. By the update rule, this means that N_3 gets doubled in most of the rounds, and is likely equal to L_1 . Thus, the update intervals constitute $\frac{L_1}{1+L_1}$ fraction of time. At the same time, in these intervals, near optimal values for N_1 are being used. Thus, we expect to see $\frac{L_1}{1+L_1}$ as a multiplicative factor in θ .

The above discussion and Theorem 1 also state that DCP successfully adapts N_1 in order to keep $\varphi(\hat{t}_k + N_c, N_1(\hat{t}_k))$ close to $\phi(\mathbf{X}(\hat{t}_k + N_c))$ ⁸. Note that for a given \mathbf{X} finding $\tilde{N}_1(\mathbf{X})$, or equivalently, $\tilde{\phi}(\mathbf{X})$, in general, is a difficult problem. Specifically, it requires the exact knowledge of the channel state and arrival process statistics, and the structure

of algorithm A . Even when this knowledge is available, as the number of users increases, finding $\tilde{N}_1(\mathbf{X})$ demands computation over a larger number of dimensions, which becomes exponentially complex. Hence, we see that DCP dynamically solves a difficult optimization problem, without requiring the knowledge of input rates or the structure of algorithm A ⁹.

2) Comparison with Static Policies, Minmax v.s. Maxmin:

Part (b) of the theorem gives the region $\theta_\infty \Gamma$ as the fundamental lower-bound on the limiting performance of DCP. It also implicitly states that this lower-bound depends on the solution to a minmax problem. To see this, recall that by definition $\tilde{\phi}(\mathbf{X})$ is the maximum of $\phi(\mathbf{X}, N_1)$ over all choices for N_1 . Thus, we have that

$$\theta_\infty = \inf_{\|\mathbf{X}\|=1} \max_{N_1 \in \mathcal{N}_1} \frac{\phi(\mathbf{X}, N_1)}{\chi(\mathbf{X})}.$$

Now, consider a *static* policy that assumes a fixed value for N_1 . This policy partitions the time axis into a set of frames each consisting of N_1 timeslots, with the i th frame starting at time $(i-1)N_1$. The static policy, in the beginning of each frame, e.g., the i th frame, provides algorithm A with vectors $\mathbf{X}((i-1)N_1)$ and $\mathbf{s}((i-1)N_1)$. Algorithm A uses these vectors as input, and after spending N_1 timeslots, returns a schedule vector as the output. This output vector is then used to schedule users in the next following frame.

It is not difficult to show that the above static policy stabilizes the network for all rates interior to $\theta_{N_1}^s \Gamma$, where

$$\theta_{N_1}^s = \inf_{\|\mathbf{X}\|=1} \frac{\phi(\mathbf{X}, N_1)}{\chi(\mathbf{X})}.$$

Thus, the best static policy, in terms of the region $\theta_{N_1}^s \Gamma$, is the one that maximizes $\theta_{N_1}^s$. Let θ_o^s be the maximum value. We have that

$$\theta_o^s = \max_{N_1 \in \mathcal{N}_1} \inf_{\|\mathbf{X}\|=1} \frac{\phi(\mathbf{X}, N_1)}{\chi(\mathbf{X})}.$$

Therefore, the best static policy corresponds to a maxmin problem. Considering the definition of θ_∞ and θ_o^s , and that the minmax of a function is always larger than or equal to the maxmin, we have that $\theta_o^s \Gamma \subseteq \theta_\infty \Gamma$. More generally, using the above definitions and a simple drift analysis, we can show that the stability region of static policies is not larger than the limiting stability region of DCP.

3) Tightness of θ_∞ and θ_o^s :

Note that parts (a) and (b) of the theorem do not exclude the possibility of networks being stable under DCP for rates outside of $\theta \Gamma$ or $\theta_\infty \Gamma$. Part (c) of the theorem, on the other hand, complements parts (a) and (b), and shows that for some networks the region $\theta_\infty \Gamma$ is indeed the largest *scaled* version of Γ that can be stably supported under DCP. This for instance may happen when the channel state is statistically symmetric with respect to users as the ones in Section VI. Proof of part (c) of the theorem provides conditions for cases that lead to the maximal stability of the region $\theta_\infty \Gamma$, and in particular, shows that the symmetric examples in Section VI meet such conditions. Note that the

⁸This statement is in fact a direct result of Lemma 4.
⁹DCP also does not require the exact knowledge of channel state statistics. However, a practical implementation of DCP requires N_c to be related to the convergence-rate of channel process to its steady state.

same discussion also applies to $\theta_{N_1}^s$ and the stability region of static policies. We therefore have θ_∞ and θ_o^s both as tight measures, stating that for some networks, including the ones in the next section, DCP can increase throughput efficiency of static policies by a factor of $\frac{\theta_\infty - \theta_o^s}{\theta_o^s}$.

4) *Delay*: Note that getting close to the boundary of $\theta_\infty \Gamma$ increases delay. This follows from part (b) of the theorem stating that for input rates close to the boundary, L_1 and N_c should be large. These choices, as expected, increase the length of test and update intervals, which can potentially be large intervals of sub-optimal transmissions in terms of the value used for N_1 . This in turn makes data wait in queues before transmission, thus increasing the delay.

5) *Distributed Implementation*: Assuming algorithm A is decentralized [7][5][6][9], DCP can be implemented in a distributed manner with low overhead. This is possible since consistent implementation of DCP in all nodes requires updates of only queue backlog and nodes' time-average of backlog-rate product, and such updates are needed only over long time intervals.

More specifically, two conditions are required to be met for distributed implementation. First, nodes should generate the same sequence of random candidates for N_1 over time, which can be met by assuming the same number generator is employed by all nodes. Second, nodes should have the knowledge of backlog-rate product in the test and its preceding update interval in order to individually and consistently apply the update rule.

The second condition can also be met, for instance, by requiring each source node perform the following. Every node, e.g., the i th node, records its own backlog, $X_i(t)$, only at the beginning of the test and update intervals. During these intervals, the i th node also computes its own *individual* time-average of backlog-rate product $X_i D_i$. Here, we assume the time-averages in the test intervals are computed up to the last N_d timeslots, where $N_d \ll N_c$. Then, once an update interval ends, the i th node has all the duration of a test interval, consisting of N_c timeslots, to send all the other nodes X_i and time-average of $X_i D_i$ for that update interval. Similarly, when the last N_d timeslots in a test interval are reached, the i th node starts sending all the other nodes X_i and time-average of $X_i D_i$ of that test interval, hence, having N_d timeslots for communication. Since for each interval, data of each node, backlog in the beginning of the associated interval and the time-average, consists of at most a few bytes, we see that the overhead can be made arbitrarily small by choosing N_c and N_d large. At the same time, we can make the ratio $\frac{N_d}{N_c}$ sufficiently small, by choosing N_c large, to ensure that not consideration of the last N_d timeslots in the test intervals has little impact on the stability region.

VI. CASE STUDIES

In this section, we present two examples that provide further insight into our analytical results and the performance of DCP. To be able to compare the simulation results with analytical ones, we consider a small network consisting of two data flows in the downlink of a wireless LAN or a cellular network. In this case, $\mathbf{s}(t)$ is the vector of channel gains, and we assume

the schedule vector is the power allocation vector, i.e., $\mathbf{I} = \mathbf{P} = (p_1, p_2)$, with constraint

$$p_1 + p_2 = P_t,$$

where P_t is total power budget. Assuming super-position coding is used in the downlink, if $s_1(t) < s_2(t)$, then [33]

$$D_1(\mathbf{s}(t), \mathbf{P}) = \log \left(1 + \frac{p_1 |s_1|^2}{p_2 |s_1|^2 + n_0} \right),$$

and

$$D_2(\mathbf{s}(t), \mathbf{P}) = \log \left(1 + \frac{p_2 |s_2|^2}{n_0} \right).$$

If $s_1(t) \geq s_2(t)$, we obtain similar expressions for user rates by swapping the role of one user for another.

For illustration purposes, we assume that algorithm A in every step, i.e., during each timeslot, reduces the gap to the optimal backlog-rate product. Specifically, if the initial gap corresponding to the initial power vector $\mathbf{I}^{(0)}$, assumed to be chosen randomly, is Δ_0 , then after i steps the gap is decreased to Δ_i , where

$$\begin{aligned} \Delta_i &= \mathbf{XD}^*(\mathbf{X}, \mathbf{s}) - \mathbf{XD}(\mathbf{s}, \mathbf{I}^{(n)}) \\ &= \frac{1}{\beta^i} (\mathbf{XD}^*(\mathbf{X}, \mathbf{s}) - \mathbf{XD}(\mathbf{s}, \mathbf{I}^{(0)})) = \frac{\Delta_0}{\beta^i}, \end{aligned}$$

where $\beta > 1$. This case corresponds to $g(n) = (1 - \zeta^n)$ with $\zeta = \frac{1}{\beta}$, where $g(n)$ is introduced in Section IV-A.

Having specified rates and algorithm A , as the first example, we assume that the channel state is Markovian with two possible state vectors, namely, $\mathbf{s}_1 = (1, 5)$ and $\mathbf{s}_2 = (5, 1)$, where the channel vector in each transition takes a different state with probability $p_t = 0.3$. For this case, we set $\alpha = 0.06$, $N_c = 12000$, $L_1 = 32$, $\beta = 1.7$, $\mathcal{N}_1 = \{N_1 : 1 \leq N_1 \leq 6\}$, $n_0 = 10$, and $p_t = 50$. To study the stability region, we consider the rate vector $\mathbf{a} = (2.4181, 2.4181)$ which belongs to the boundary of Γ corresponding to this example. We then assume the arrival vector is $\gamma \mathbf{a}$, where γ is the load factor, and varies from 0.84 to 0.92. Fig. 2 depicts the resulting average queue sizes. For loads larger than 0.93, the queue sizes increase with time implying network instability. The range selected for γ is motivated by noting that $\theta_\infty = 0.9447$, which is computed numerically. Considering the growth of average queue sizes in Fig. 2, we therefore see that for this example θ_∞ is indeed an upper bound for capacity region scaling. In fact, part (c) of Theorem 1 applies to this example, and any rate of the form $(\theta_\infty + \epsilon)\mathbf{a}$, $\epsilon > 0$, makes the network unstable.

As for the second example, we increase the number of states to six corresponding to the following state vectors:

$$\begin{aligned} \mathbf{s}_1 &= (1, 5), \quad \mathbf{s}_2 = (5, 1), \\ \mathbf{s}_3 &= (1, 2), \quad \mathbf{s}_4 = (2, 1), \\ \mathbf{s}_5 &= (2, 5), \quad \mathbf{s}_6 = (5, 2), \end{aligned}$$

and having the following symmetric transition matrix:

$$\mathbf{T}_m = \begin{pmatrix} 0.3 & 0.1 & 0.2 & 0.1 & 0.2 & 0.1 \\ 0.1 & 0.3 & 0.1 & 0.2 & 0.1 & 0.2 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0.1 & 0.2 & 0.1 & 0.2 & 0.1 & 0.3 \end{pmatrix}. \quad (9)$$

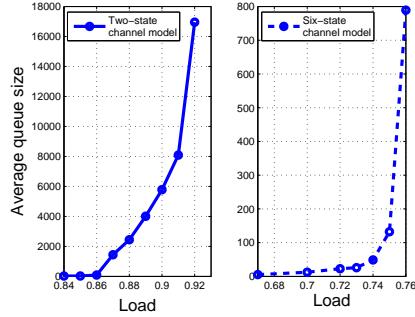


Fig. 2. Average queue size as a function of load factor.

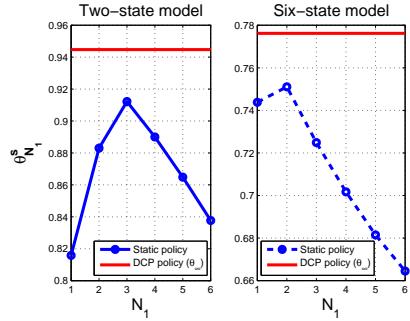


Fig. 3. Comparison of capacity region scaling for DCP and static policies.

For this case, we keep the same N_c , L_1 , and \mathcal{N}_1 , but assume $\alpha = 0.02$, $\beta = 1.5$, $n_0 = 50$, and $p_t = 10$. Similar to the previous example, to vary arrival rate vector, we consider the rate vector $\mathbf{a} = (0.6952, 0.6952)$ which belongs to the boundary of Γ associated with this example. Then, the arrival vector is assumed to be $\gamma\mathbf{a}$, where the load factor γ varies from 0.67 to 0.76. The resulting average queue sizes are also shown in Fig. 2. In this case, for load factors larger than 0.76, the queue sizes increase with time, suggesting network instability. This result is consistent with our analytical results since the numerically computed value of θ_∞ is 0.7762. Note that part (c) of Theorem 1 also applies to this example, and any rate of the form $(\theta_\infty + \epsilon)\mathbf{a}$, $\epsilon > 0$, makes the network unstable.

Finally, in Fig. 3, for the two examples, we have shown $\theta_{N_1}^s$ as a function of N_1 , and also shown the value of θ_∞ for DCP. As expected and the figure suggests, since DCP adapts N_1 according to queue backlog, it outperforms the best static policy. We also see that the optimal stationary policy for the first example is the one with $N_1 = 3$ and $\theta_o^s = 0.9122$, and for the second example is the one with $N_1 = 2$ and $\theta_o^s = 0.7511$. Note that characterization of the best static policy requires computation of $\tilde{\phi}(\mathbf{X})$, which, as briefly discussed in Section V-D1, can be computationally intensive. From the figure, we also observe that the performance of a suboptimal static policy can be substantially less than DCP if the static policy does not assume a proper value for N_1 .

VII. CONCLUSION

In this paper, to improve the stable throughput region in practical network setups, we have considered the problem of scheduling in time-varying networks from a new perspective. Specifically, in contrast to previous research which assumes the search-time to find schedule vectors is negligible, we

have considered this time, based on which we modeled the time-efficiency of sub-optimal algorithms. Inspired by this modeling, we have proposed a dynamic control policy that dynamically but in a large time-scale tunes the time given to an available sub-optimal algorithm according to queue backlog and channel correlation. Remarkably, this policy does not require knowledge of input rates or the structure of available sub-optimal algorithms, nor it requires exact statistics of the channel process. We have shown that this policy can be implemented in a distributed manner with low overhead. In addition, we have analyzed the throughput stability region of the proposed policy and shown that its throughput region is at least as large as the one for any other, including the optimal, static policy. We believe that study and design of similar policies opens a new dimension in the design of scheduling policies, and in parallel to the efforts to improve the performance of sub-optimal algorithms, can help boost the throughput performance to the capacity limit.

APPENDIX

PROOF OF THEOREM 1

Proof of part (a):

The proof of part (a) consists of two main parts. First, using several lemmas, we obtain a negative drift with a random number of steps. In the second part, we use the negative drift analysis to show that the return time to a bounded region has a finite expected value, and conforms to the properties required for network stability, according to the definition given in Section V-A1.

We start by noting that $\theta \leq 1$, and since \mathbf{a} is strictly inside $\theta\Gamma$, there must be some non-negative constants $\beta_{s,I}$ with the property that for all $s \in \mathcal{S}$

$$\sum_{I \in \mathcal{I}} \beta_{s,I} < \theta \leq 1, \quad (10)$$

such that

$$\mathbf{a} = \sum_{s \in \mathcal{S}} \pi(s) \sum_{I \in \mathcal{I}} \beta_{s,I} \mathbf{D}_{s,I}. \quad (11)$$

Considering (10), we can define positive ξ' as

$$\xi' = \theta - \max_{s \in \mathcal{S}} \sum_{I \in \mathcal{I}} \beta_{s,I}.$$

Since $\xi' > 0$, by the definition of θ , for $\|\mathbf{X}_t\| \neq 0$, we have that

$$\frac{R_\infty(\tilde{\phi}(t) - \alpha - 3\theta_\varphi)}{\chi(\mathbf{X}_t)} - \max_{s \in \mathcal{S}} \sum_{I \in \mathcal{I}} \beta_{s,I} > \xi' > 0. \quad (12)$$

To proceed with the proof, associated with a given time t , we define a sequence of random variables $\{\tau_i\}_{i=-1}^\infty$, where τ_{-1} and τ_0 denote the number of timeslots to the last timeslot of the previous and the current scheduling round, respectively, and τ_i , $i \geq 1$, is the number of timeslots to the last timeslot of the i th subsequent scheduling round. Let \mathcal{H}_t denote the past history of the system up to and including time t . Thus, given \mathcal{H}_t , the value of \mathbf{X}_t is known. Let $f(\cdot)$ be defined as

$$f(\mathbf{X}) = \|\mathbf{X}\|^2,$$

Considering a $\tau_K + 1$ -step drift with function $f(\cdot)$, we can write

$$\begin{aligned}\Delta(\tau_K + 1) &= \mathbb{E}[f(\mathbf{X}_{t+\tau_K+1}) - f(\mathbf{X}_t) | \mathcal{H}_t] \\ &= \mathbb{E}\left[\sum_{k=0}^{\tau_K} f(\mathbf{X}_{t+k+1}) - f(\mathbf{X}_{t+k}) | \mathcal{H}_t\right] \\ &= \mathbb{E}\left[\sum_{k=0}^{\tau_K} (\mathbf{X}_{t+k+1} - \mathbf{X}_{t+k})(\mathbf{X}_{t+k+1} - \mathbf{X}_{t+k})^* | \mathcal{H}_t\right].\end{aligned}$$

Using the fact that arrivals and departures are bounded, after performing some preliminary steps, we can show that

$$\begin{aligned}\Delta(\tau_K + 1) &\leq \mathbb{E}\left[(\tau_K + 1)C_1 + (\tau_K + 1)^2C_2\right. \\ &\quad \left. + 2\sum_{k=0}^{\tau_K} (\mathbf{X}_t \mathbf{A}_{t+k} - \mathbf{X}_{t+k} \mathbf{D}_{t+k})^* | \mathcal{H}_t\right],\end{aligned}$$

for appropriate constants C_1 and C_2 . Since $\mathbf{X}_{t+k} \mathbf{D}_{t+k} \geq 0$, we have

$$\begin{aligned}\Delta(\tau_K + 1) &\leq \mathbb{E}\left[(\tau_K + 1)C_1 + (\tau_K + 1)^2C_2\right. \\ &\quad \left. + 2\sum_{k=0}^{\tau_K} (\mathbf{X}_t \mathbf{A}_{t+k} - \mathbf{X}_t \mathbf{a})\right. \\ &\quad \left. + 2\sum_{k=0}^{\tau_K} (\mathbf{X}_t \mathbf{a} - \mathbf{X}_{t+k} \mathbf{D}_{t+k}^*)\right. \\ &\quad \left. + 2\sum_{k=0}^{\tau_K} \mathbf{X}_{t+k} \mathbf{D}_{t+k}^*\right. \\ &\quad \left. - 2\sum_{k=\tau_0+1}^{\tau_K} \mathbf{X}_{t+k} \mathbf{D}_{t+k} | \mathcal{H}_t\right],\end{aligned}$$

where $\mathbf{D}_{t+k}^* = \mathbf{D}^*(\mathbf{X}(t+k), \mathbf{s}(t+k))$. In the following, we derive an upper bound for $\Delta(\tau_K + 1)$.

As mentioned in Section III-A, arrivals are i.i.d with mean vector \mathbf{a} . We can therefore apply the same method used to prove Lemma 1 to obtain

$$\mathbb{E}\left[\left\|\sum_{k=0}^{\tau_K} \mathbf{A}_{t+k} - (\tau_K + 1)\mathbf{a}\right\| | \mathcal{H}_t\right] \leq \epsilon \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t],$$

where $\epsilon > 0$, and can be made arbitrarily small by choosing a sufficiently large K .

Using the above inequality, Lemma 2, Lemma 3, and Lemma 4, all with the same choice for ϵ , we can show that

$$\begin{aligned}\Delta(\tau_K + 1) &\leq \mathbb{E}\left[(\tau_K + 1)\|\mathbf{X}_t\| \chi(\mathbf{X}_t)\right. \\ &\quad \left. \left(\epsilon_1 + 2\left(\max_{\mathbf{s} \in \mathcal{S}} \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}} - \frac{R_\infty(\tilde{\phi}(t) - \alpha - 3\theta_\varphi)}{\chi(\mathbf{X}_t)}\right)\right) | \mathcal{H}_t\right],\end{aligned}\tag{13}$$

where

$$\epsilon_1 = \frac{1}{\chi(\mathbf{X}_t)} \left(\frac{C_1}{\|\mathbf{X}_t\|} + \frac{C_2(\tau_K + 1)}{\|\mathbf{X}_t\|} + 8\epsilon \right).\tag{14}$$

Note that according to the lemmas, ϵ can take any given positive real number if K and $\|\mathbf{X}_t\|$ are sufficiently large.

Similarly, ϵ_1 can assume any given positive value. To see this, first note that since $\mathbf{a} \in \theta\Gamma$, we have $\mathbf{a} \in \Gamma$. Thus, for any user, e.g. the i th user, for which $\mathbf{a}_i > 0$, there has to be a state \mathbf{s} and a schedule \mathbf{I} satisfying

$$\pi(\mathbf{s}) \mathbf{D}(\mathbf{s}, \mathbf{I})_i > 0,$$

where $\mathbf{D}(\mathbf{s}, \mathbf{I})_i$ is the i th element of vector $\mathbf{D}(\mathbf{s}, \mathbf{I})$. Otherwise, \mathbf{a}_i should be zero, contradicting the assumption. Therefore, assuming $\mathbf{a} \neq 0$, we can define positive v as

$$v = \min_{i \in N} \max_{\mathbf{s}, \mathbf{I}} \pi(\mathbf{s}) \mathbf{D}(\mathbf{s}, \mathbf{I})_i > 0.$$

Thus,

$$\mathbb{E}[\mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s})] \geq v \max_{i \in N} \mathbf{X}_t(i) \geq \frac{v}{\sqrt{N}} \|\mathbf{X}_t\|. \tag{15}$$

This implies that for all nonzero $\mathbf{X} \in \mathbb{R}^N$

$$\chi(\mathbf{X}) \geq \frac{v}{\sqrt{N}}. \tag{16}$$

On the other hand, since departure rates are bounded above by D_{max} , we have

$$\chi(\mathbf{X}) \leq \sqrt{N} D_{max}. \tag{17}$$

Now consider any positive ϵ_2 , and suppose K is sufficiently large such that for large $\|\mathbf{X}_t\|$ we have

$$\frac{8\epsilon}{\chi(\mathbf{X}_t)} \leq \frac{8\sqrt{N}\epsilon}{v} < \frac{\epsilon_2}{3},$$

where the first inequality follows from (16). This upper-bounds the third term in ϵ_1 . Since for any K , and in particular, the chosen one, we have $\tau_K + 1 \leq (K + 1)(1 + L_1)N_c$, we see that if $\|\mathbf{X}_t\|$ is appropriately large, the first and second terms in ϵ_1 can also be less than $\frac{\epsilon_2}{3}$. Thus, for any given positive ϵ_2 , we can find an appropriately large K such that for sufficiently large $\|\mathbf{X}_t\|$, (13) holds with $\epsilon_1 < \epsilon_2$.

Suppose K is sufficiently large, and $\|\mathbf{X}_t\| > M_K$ for appropriately large M_K such that $\epsilon_1 < \xi'$. We can use (13) and (12) to show that

$$\Delta(\tau_K + 1) < -\mathbb{E}[\xi' \|\mathbf{X}_t\| (\tau_K + 1) \chi(\mathbf{X}_t) | \mathcal{H}_t].$$

This inequality and (16) further imply that

$$\Delta(\tau_K + 1) < -\mathbb{E}[\xi(\tau_K + 1) \|\mathbf{X}_t\| | \mathcal{H}_t], \tag{18}$$

where $\xi = \frac{v}{\sqrt{N}} \xi' > 0$. We, therefore, have obtained the negative drift expression, completing the first part of the proof.

Note that in above τ_K is a random variable, and in fact, is a stopping time with respect to the filtration $\mathcal{H} = \{\mathcal{H}_t\}_{t=0}^\infty$. This means that we have obtained a drift expression that is based on a random number of steps. Proofs of stability in the literature, however, are often based on a negative drift with a fixed number of steps. This contrast has motivated us to adopt an interesting method recently developed in [10]. This method is general since it can be applied in both cases, and also leads to an intuitive notation of stability. However, it has been originally developed for Markov chains. Therefore, as well as using less technical notations, in what follows, we apply

minor modifications to the method so that it is appropriate in our context.

We now, in the second part of the proof, use the negative drift, and prove that the expected value of the return time to some bounded region is finite in a manner that renders network stable. Let \mathcal{C} denote the bounded region, and be defined as

$$\mathcal{C} = \{X \in \mathbb{R}^N, \|X\| \leq M_K\}.$$

Associated with \mathcal{C} , we define $\sigma_{\mathcal{C}}$ to be the number timeslots after which the process $\{\mathbf{X}_{t+i}\}_{i=0}^{\infty}$ enters \mathcal{C} , i.e.,

$$\sigma_{\mathcal{C}} = \inf\{i \geq 0 : \mathbf{X}_{t+i} \in \mathcal{C}\}.$$

Similarly, we let $\tau_{\mathcal{C}}$ be

$$\tau_{\mathcal{C}} = \inf\{i \geq 1 : \mathbf{X}_{t+i} \in \mathcal{C}\}.$$

Therefore, $\tau_{\mathcal{C}}$, in contrast $\sigma_{\mathcal{C}}$, characterizes the first time that the process $\{\mathbf{X}_{t+i}\}_{i=1}^{\infty}$ returns to \mathcal{C} .

Back to the drift expression in (18), let η be a random variable defined by

$$\eta = \xi(\tau_K + 1)\|\mathbf{X}_t\|.$$

We obtain, for K sufficiently large,

$$\mathbb{E}[f(\mathbf{X}_{t+\tau_K+1}) + \eta | \mathcal{H}_t] \leq f(\mathbf{X}_t), \quad (19)$$

provided that $\|\mathbf{X}_t\| > M_K$. Let $\eta_0 = \eta$, and $\tau_{K,0} = \tau_K$, where η and τ_K are random variables defined by considering time t . We now consider time $t_K^{(1)} = t + \tau_{K,0} + 1$. For this particular time, we can define another pair $\tau_{K,1}$ and η_1 and such that if $\|\mathbf{X}_{t_K^{(1)}}\| > M_K$, then

$$\mathbb{E}[f(\mathbf{X}_{t_K^{(1)} + \tau_{K,1} + 1}) + \eta_1 | \mathcal{H}_{t_K^{(1)}}] \leq f(\mathbf{X}_{t_K^{(1)}}),$$

where $\tau_{K,1}$ is the number of timeslots from time $t_K^{(1)}$ to the last timeslot of the K th subsequent scheduling round, and

$$\eta_1 = \xi(\tau_{K,1} + 1)\|\mathbf{X}_{t_K^{(1)}}\|.$$

Note that the definition of $\tau_{K,1}$ and η_1 is independent of whether the previous inequality holds.

We can continue this process by considering the drift criteria for time $t_K^{(i)} = t_K^{(i-1)} + \tau_{K,i-1} + 1$, and defining random variables $\tau_{K,i}$ and η_i . The random variables $\tau_{K,i}$ and η_i have a similar definition as $\tau_{K,1}$ and η_1 , respectively, except that they are associated with time $t_K^{(i)}$. Using these definitions, we can define $t_K^{(i)}$ more precisely by

$$\begin{aligned} t_K^{(0)} &= t, \\ t_K^{(i)} &= t_K^{(i-1)} + (\tau_{K,i-1} + 1) = t + \sum_{j=0}^{i-1} (\tau_{K,j} + 1). \end{aligned}$$

Note that $t_K^{(i)}$ is a stopping time with respect to \mathcal{H} . Using $t_K^{(i)}$, we set

$$\bar{\mathbf{X}}_i = \mathbf{X}_{t_K^{(i)}}, \quad i \geq 0, \quad (20)$$

and define \mathcal{H}^{τ} as the filtration given by $\mathcal{H}^{\tau} = \{\mathcal{H}_{t_K^{(i)}}\}_{i=0}^{\infty}$. In addition, associated with η_i , which is given by

$$\eta_i = \xi(\tau_{K,i} + 1)\|\mathbf{X}_{t_K^{(i)}}\|,$$

we define $\eta^{(i)}$ as

$$\eta^{(0)} = 0, \quad \eta^{(i)} = \sum_{j=0}^{i-1} \eta_j. \quad (21)$$

We also define ν as

$$\nu = \inf\{i \geq 0 : t_K^{(i)} \geq t + \sigma_{\mathcal{C}}\}, \quad (22)$$

which is a stopping time with respect to \mathcal{H}^{τ} . Intuitively, ν marks the first time $t_K^{(i)}$ at or before which the process $\{\mathbf{X}_{t+i}\}_{i=0}^{\infty}$ enters \mathcal{C} . We finish the chain of definitions by introducing the sequence $\{Z_i\}_{i=0}^{\infty}$, where

$$Z_i = f(\bar{\mathbf{X}}_i) + \eta^{(i)}. \quad (23)$$

For $i < \nu$, using (21), we have

$$\begin{aligned} \mathbb{E}[Z_{i+1} | \mathcal{H}_{t_K^{(i)}}] &= \mathbb{E}[f(\bar{\mathbf{X}}_{i+1}) + \eta_i | \mathcal{H}_{t_K^{(i)}}] + \eta^{(i)} \\ &\leq f(\bar{\mathbf{X}}_i) + \eta^{(i)} = Z_i, \end{aligned} \quad (24)$$

where the first equality follows from the fact that $\eta^{(i)}$ is completely determined given $\mathcal{H}_{t_K^{(i)}}$, and the inequality is simply an immediate result of (19) and the assumption $i < \nu$. To simplify the notation, let $\nu \wedge i$ denote

$$\nu \wedge i = \min(\nu, i).$$

It now follows directly from (24) that the sequence $\{Z_{\nu \wedge i}\}_{i=0}^{\infty}$ is an \mathcal{H}^{τ} -supermartingale. Since $f(\cdot)$ is non-negative, we have

$$\mathbb{E}[\eta^{(\nu \wedge i)} | \mathcal{H}_t] \leq \mathbb{E}[Z_{\nu \wedge i} | \mathcal{H}_t].$$

But $\mathcal{H}_t = \mathcal{H}_{t_K^{(0)}}$, and $\{Z_{\nu \wedge i}\}_{i=0}^{\infty}$ is a supermartingale. Hence,

$$\begin{aligned} \mathbb{E}[Z_{\nu \wedge i} | \mathcal{H}_t] &= \mathbb{E}[Z_{\nu \wedge i} | \mathcal{H}_{t_K^{(0)}}] \\ &\leq Z_0 = f(\mathbf{X}_t). \end{aligned}$$

Considering the last two inequalities, we obtain

$$\mathbb{E}[\eta^{(\nu \wedge i)} | \mathcal{H}_t] \leq f(\mathbf{X}_t). \quad (25)$$

In addition, using the definition of $\eta^{(i)}$ and η_j while assuming $M_K > 1$, it is easy to see that

$$\begin{aligned} \eta^{(\nu \wedge i)} &= \sum_{j=0}^{i-1} \eta_j \mathbf{1}_{(j < \nu)} \geq \xi \sum_{j=0}^{i-1} (\tau_{K,j} + 1) \mathbf{1}_{(j < \nu)} \\ &= \xi(t_K^{(\nu \wedge i)} - t). \end{aligned} \quad (26)$$

Applying the monotone convergence theorem [29], we can take the limit in (25) and (26) as $i \rightarrow \infty$ yielding

$$\mathbb{E}[t_K^{(\nu)} - t | \mathcal{H}_t] \leq \xi^{-1} f(\mathbf{X}_t).$$

But by definition in (22), $\sigma_{\mathcal{C}} \leq t_K^{(\nu)} - t$. Thus, for $\mathbf{X}_t \notin \mathcal{C}$

$$\mathbb{E}[\sigma_{\mathcal{C}} | \mathcal{H}_t] \leq \xi^{-1} f(\mathbf{X}_t).$$

If $\mathbf{X}_t \in \mathcal{C}$, we have $\sigma_{\mathcal{C}} = 0$. Hence, we have that

$$\mathbb{E}[\sigma_{\mathcal{C}} | \mathcal{H}_t] \leq \xi^{-1} f(\mathbf{X}_t) \mathbf{1}_{\mathbf{X}_t \notin \mathcal{C}},$$

showing that the expected $\sigma_{\mathcal{C}}$ is bounded by a function of \mathbf{X}_t uniformly in the past history and t , as required. This completes the proof of part (a) of the theorem. \blacksquare

Proof of part (b): Part (b) follows directly from part (a) of the theorem as a corollary by noting that θ_φ and ϱ_φ can be made arbitrarily small by assuming a sufficiently large N_c , as stated in Property 2. This allows us to select arbitrarily small values for α . In addition, we can chose a sufficiently large value for L_1 such that for sufficiently small values for θ_φ and ϱ_φ , R_∞ is arbitrarily close to one. Considering these choices, we see that we can make θ arbitrarily close to θ_∞ , as required. \blacksquare

Proof of part (c): Since part (c) of the theorem only concerns existence of such networks for which the region $\theta_\infty\Gamma$ is maximally stable under DCP, for simplicity of exposition, we consider a network consisting of two users, i.e., two data flows. Note that our approach can be extended to more general networks with N data flows. Here, we adopt a direct method and show that with positive probability norm of the backlog vector approaches infinity. Therefore, the expected value of the return time to any bounded region becomes infinity, implying network instability. We start by introducing several definitions followed by four conditions sufficient for network instability.

Let $\bar{\mathbf{D}}$ and $\bar{\mathbf{D}}^*$ be defined by

$$\bar{\mathbf{D}}(\mathbf{X}) = \mathbb{E} \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} (\mathbf{D}_{t+i} - \mathbf{U}_{t+i}) \mid \mathbf{X}_{t+i} = \mathbf{X}, i \geq 0 \right],$$

and

$$\bar{\mathbf{D}}^*(\mathbf{X}) = \mathbb{E}_{\mathbf{s}} [\mathbf{D}^*(\mathbf{X}, \mathbf{s})],$$

where $\mathbf{D}^*(\mathbf{X}, \mathbf{s})$ is defined in (3). In addition, let \mathbf{X}_{min} be¹⁰

$$\mathbf{X}_{min} = \underset{\|\mathbf{X}\|=1}{\operatorname{arginf}} \frac{\tilde{\phi}(\mathbf{X})}{\chi(\mathbf{X})}.$$

Note that in the definition of $\bar{\mathbf{D}}$, we hypothetically assume that the backlog vector after time t is fixed and does not change. This is similar to the method used to define $\phi(\mathbf{X}, N_1)$ except that here we do not assume a fixed value for N_1 , and instead, assume DCP adapts N_1 as if the backlog vector was changing. In addition, note that by the ergodicity of the channel process $\bar{\mathbf{D}}$ does not depend on t , and moreover, by Property 1, $\bar{\mathbf{D}}$ does not depend on $\|\mathbf{X}\|$. To simplify the subsequent analysis, we also consider the following definitions:

Definition 1: For a given \mathbf{X} and a given $\epsilon > 0$, the ϵ -neighborhood of \mathbf{X} is defined by

$$\mathcal{N}(\mathbf{X}, \epsilon) = \{\mathbf{X}_1 : \|\mathbf{X}_1 - \mathbf{X}\| < \epsilon\}.$$

Definition 2: For a given \mathbf{X} with $\|\mathbf{X}\| = 1$, and a given $\epsilon > 0$, the *normalizing* region $\mathcal{R}(\mathbf{X}, \epsilon)$ is defined by

$$\mathcal{R}(\mathbf{X}, \epsilon) = \{\mathbf{X}_1 : \|\mathbf{X}_1\| \neq 0, \left\| \frac{\mathbf{X}_1}{\|\mathbf{X}_1\|} - \mathbf{X} \right\| < \epsilon\} \cup \{\mathbf{X}_1 = \mathbf{0}\}.$$

Definition 3: Consider a region \mathcal{R} and a vector \mathbf{X} inside \mathcal{R} . We define $\xi(\mathbf{X}, \mathcal{R})$ as the *supremum* of the angular deviation of the vectors in \mathcal{R} from \mathbf{X} , i.e.,

$$\xi(\mathbf{X}, \mathcal{R}) = \sup_{\mathbf{Y} \in \mathcal{R}} \arccos \left(\frac{\mathbf{X} \cdot \mathbf{Y}}{\|\mathbf{X}\| \|\mathbf{Y}\|} \right).$$

¹⁰Note that here infimum can be achieved since the functions $\tilde{\phi}(\mathbf{X})$ and $\chi(\mathbf{X})$ are continuous functions of \mathbf{X} , and the infimum is taken over a closed interval.

Now suppose the following conditions hold:

- C1) $\mathbf{X}_{min} = \gamma_1 \bar{\mathbf{D}}(\mathbf{X}_{min}) = \gamma_2 \bar{\mathbf{D}}^*(\mathbf{X}_{min})$, for some $\gamma_1, \gamma_2 > 0$.
- C2) For any $N_{1,1} \in \mathcal{N}_1$ and $N_{1,2} \in \mathcal{N}_1$ with $N_{1,1} \neq N_{1,2}$, we have $\phi(\mathbf{X}_{min}, N_{1,1}) \neq \phi(\mathbf{X}_{min}, N_{1,2})$.
- C3) For any $\beta_1 > 0$ and $\beta_2 > 0$, there exists a sufficiently small $\epsilon > 0$ such that if $\mathbf{X} \in \mathcal{R}(\mathbf{X}_{min}, \epsilon)$, then

$$\bar{\mathbf{D}}(\mathbf{X}) - \bar{\mathbf{D}}(\mathbf{X}_{min}) = \lambda_1 \bar{\mathbf{D}}(\mathbf{X}_{min}) + \lambda_2 \left(\frac{\mathbf{X}}{\|\mathbf{X}\|} - \mathbf{X}_{min} \right),$$

for some λ_1 and λ_2 satisfying $|\lambda_1| < \beta_1$ and $0 < \lambda_2 < \beta_2$.

- C4) For any $\mathbf{X} \in \mathbb{N}^N$, for some t

$$P(\mathbf{X}_t = \mathbf{X}) > 0.$$

Condition C1 may be met by assuming a statistically symmetric channel states as the ones in Section VI. Condition C2 simply requires the function $\phi(\mathbf{X}_{min}, N_1)$ to be a one-to-one function of N_1 at \mathbf{X}_{min} . Condition C3 intuitively states that the average departure rates should be a *continuous* function¹¹ of \mathbf{X} around \mathbf{X}_{min} , and in particular, when \mathbf{X} deviates from \mathbf{X}_{min} , these rates should deviate from $\bar{\mathbf{D}}(\mathbf{X}_{min})$ in a similar manner. This is in fact expected as increasing the backlog vector in one dimension should increase the expected departure rate in that dimension, which can be considered as a result of the approximation to the GMWM problem through the use of algorithm A. Note that in C3 where appropriate the vector \mathbf{X} is normalized by its norm since $\bar{\mathbf{D}}(\mathbf{X})$ does not depend on $\|\mathbf{X}\|$. Finally, C4 simply requires the process $\{\mathbf{X}_t\}$ to be able to reach all vectors in \mathbb{N}^N , although what we need for the proof is a relaxed version of this assumption. Using the numerical results for $\phi(\mathbf{X}_{min}, N_1)$ and $\bar{\mathbf{D}}^*(\mathbf{X}_{min})$, and the symmetry of channel states, it is easy to verify that the conditions C2-C4 also hold for the examples in Section VI. Therefore, there are examples for which the conditions C1-C4 hold. Next, we show that these conditions are sufficient for network instability, completing the proof of part (c).

First, note that $\bar{\mathbf{D}}^*(\mathbf{X}_{min}) \in \Gamma$, which directly follows from the definition of $\bar{\mathbf{D}}^*(\mathbf{X}_{min})$ and Γ . Second, the rate $\bar{\mathbf{D}}^*(\mathbf{X}_{min})$ belongs to the boundary of Γ , otherwise we could find another vector \mathbf{D} inside Γ and within a small neighborhood of $\bar{\mathbf{D}}^*(\mathbf{X}_{min})$ with larger backlog-rate product, in contradiction with the definition of $\bar{\mathbf{D}}^*(\mathbf{X}_{min})$. Hence, we see that the rate $\theta_\infty \bar{\mathbf{D}}^*(\mathbf{X}_{min})$ belongs to the boundary of $\theta_\infty \Gamma$. Third, we can see that by the definition of θ_∞ and \mathbf{X}_{min}

$$\theta_\infty = \frac{\tilde{\phi}(\mathbf{X}_{min})}{\chi(\mathbf{X}_{min})} \geq \frac{\mathbf{X}_{min} \bar{\mathbf{D}}(\mathbf{X}_{min})}{\chi(\mathbf{X}_{min})}.$$

This is because DCP may use sub-optimal values for N_1 , which by C2 make $\mathbf{X}_{min} \bar{\mathbf{D}}$ less than $\tilde{\phi}(\mathbf{X}_{min})$ when N_c is large. Using C1 and the above inequality, we have

$$\theta_\infty \geq \frac{\|\bar{\mathbf{D}}(\mathbf{X}_{min})\|}{\|\bar{\mathbf{D}}^*(\mathbf{X}_{min})\|},$$

¹¹As opposed to traditional definitions which usually use $\mathcal{N}(\mathbf{X}, \epsilon)$ to define continuity, here, the region $\mathcal{R}(\mathbf{X}, \epsilon)$ is used to characterize continuity.

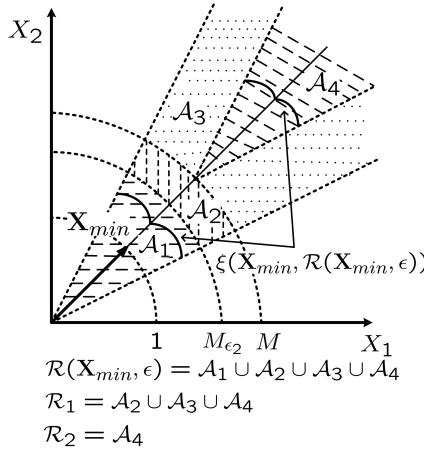


Fig. 4. Illustration of regions $\mathcal{R}(\mathbf{X}_{min}, \epsilon)$, \mathcal{R}_1 , and \mathcal{R}_2 .

which implies that

$$\bar{\mathbf{D}}(\mathbf{X}_{min}) \leq \theta_\infty \bar{\mathbf{D}}^*(\mathbf{X}),$$

where the inequality is component-wise. Without loss of generality, we assume that

$$\bar{\mathbf{D}}(\mathbf{X}_{min}) = \theta_\infty \bar{\mathbf{D}}^*(\mathbf{X}). \quad (27)$$

Let the input rate be

$$\mathbf{a} = (\theta_\infty + \varsigma) \bar{\mathbf{D}}^*(\mathbf{X}), \quad (28)$$

for some $\varsigma > 0$, which is clearly outside of the region $\theta_\infty \Gamma$. Let $\Delta_{t,n}$ be the drift vector defined by

$$\Delta_{t,n} = \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{A}_{t+i} - \frac{1}{n} \sum_{i=0}^{n-1} (\mathbf{D}_{t+i} - \mathbf{U}_{t+i}).$$

In addition, let

$$\bar{\Delta}_{\mathbf{X}} = \mathbf{a} - \bar{\mathbf{D}}(\mathbf{X}).$$

Note that $\bar{\Delta}_{\mathbf{X}}$ does not depend on $\|\mathbf{X}\|$ since $\bar{\mathbf{D}}(\mathbf{X})$ has the same property as pointed out earlier. Suppose for a given ϵ_1 , the values for β_1 and β_2 are chosen such that by C3 if $\mathbf{X}_t \in \mathcal{R}(\mathbf{X}_{min}, \epsilon)$, for appropriately small ϵ , then the following holds

$$\|\bar{\Delta}_{\mathbf{X}_t} - \bar{\Delta}_{\mathbf{X}_{min}}\| < \epsilon_1.$$

Using Assumption 1, condition C2, channel ergodicity as stated in Section III-B, and that arrivals are i.i.d, it is not hard to see that¹² when ϵ is sufficiently small, for any positive ϵ_2 and $0 < \varsigma \leq 1$, we can first chose n large and then M_{ϵ_2} sufficiently large, and define the region \mathcal{R}_1 as

$$\mathcal{R}_1 = \{\mathbf{X} : \|\mathbf{X}\| \geq M_{\epsilon_2}, \mathbf{X} \in \mathcal{R}(\mathbf{X}_{min}, \epsilon)\}$$

such that

$$P(\|\Delta_{t,n} - \bar{\Delta}_{\mathbf{X}_t}\| < \epsilon_2 | \mathcal{H}_t, \mathbf{X}_t \in \mathcal{R}_1) > (1 - \varsigma), \quad (29)$$

where in above

$$\bar{\Delta}_{\mathbf{X}_t} = \frac{(\varsigma - \lambda_1 \theta_\infty)}{\theta_\infty + \varsigma} \mathbf{a} - \lambda_2 \left(\frac{\mathbf{X}_t}{\|\mathbf{X}_t\|} - \mathbf{X}_{min} \right). \quad (30)$$

¹²A similar discussion similar to the one for Property 2 applies here.

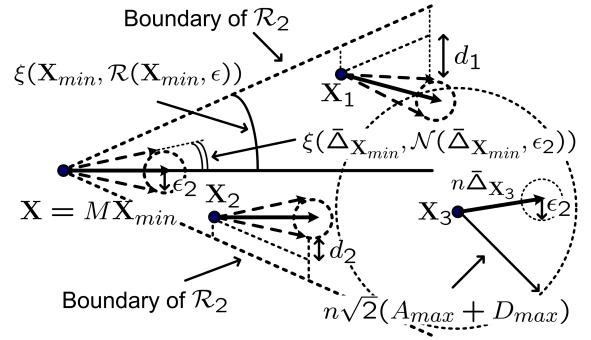


Fig. 5. Examples where $\mathbf{X}_{t+ni} \in \mathcal{R}_2$ explaining cases where $\mathcal{A}_{t+ni,n} = 0$ as in the points \mathbf{X} , \mathbf{X}_1 , and \mathbf{X}_2 , and the cases where $\mathcal{A}_{t+ni,n} = 1$ as in the point \mathbf{X}_3 . In this figure, the region \mathcal{R}_2 is rotated clockwise.

The above equality is obtained by using condition C3, equality (27), and considering that the input rate is given by (28). In particular, we have that

$$\bar{\Delta}_{\mathbf{X}_{min}} = \frac{\varsigma}{\theta_\infty + \varsigma} \mathbf{a}.$$

Since ϵ_2 can be made arbitrarily small by choosing sufficiently large n and M_{ϵ_2} , we assume that for all $\mathbf{X} \in \mathcal{R}_1$

$$\xi(\bar{\Delta}_{\mathbf{X}}, \mathcal{N}(\bar{\Delta}_{\mathbf{X}}, \epsilon_2)) < \frac{\xi(\mathbf{X}_{min}, \mathcal{R}(\mathbf{X}_{min}, \epsilon))}{2}. \quad (31)$$

Hence, according to (29) and (31), for $\mathbf{X}_t \in \mathcal{R}_1$, with probability larger than $(1 - \varsigma)$ the drift $\Delta_{t,n}$ is close to $\bar{\Delta}_{\mathbf{X}_t}$ with a supremum angular deviation that is half of the supremum angular deviation of \mathbf{X} 's in $\mathcal{R}(\mathbf{X}_{min}, \epsilon)$ from \mathbf{X}_{min} .

To continue, let the region \mathcal{R}_2 be defined as

$$\mathcal{R}_2 = \{\mathbf{X} : \mathbf{X} - M\mathbf{X}_{min} \in \mathcal{R}(\mathbf{X}_{min}, \epsilon)\}, \quad (32)$$

for some $M \geq M_{\epsilon_2}$. This region is a shifted version of $\mathcal{R}(\mathbf{X}_{min}, \epsilon)$ with the origin shifted to $M\mathbf{X}_{min}$, and therefore, $\mathcal{R}_2 \subset \mathcal{R}_1$. Fig. 4 provides a graphical demonstration of regions $\mathcal{R}(\mathbf{X}_{min}, \epsilon)$, \mathcal{R}_1 , and \mathcal{R}_2 . In the figure, the vector \mathbf{X}_{min} is shown by a unit arrow-vector. Now we are in a position to show that starting at $\mathbf{X}_t = M\mathbf{X}_{min}$, for some appropriately chosen M , with positive probability $\{\mathbf{X}_{t+i}, i \geq 0\}$ stays in \mathcal{R}_2 with ever growing norm.

Consider the sequence $\{\mathbf{X}_{t+ni}\}_{i=0}^\infty$ with $\mathbf{X}_t = M\mathbf{X}_{min}$. Recall that n is chosen sufficiently large according to the value of ϵ_2 . Let $\mathcal{A}_{t+ni,n}$ be a r.v. defined by

$$\mathcal{A}_{t+ni,n} = \begin{cases} 1 & \text{if } \|\Delta_{t+ni,n} - \bar{\Delta}_{\mathbf{X}_{t+ni}}\| \leq \epsilon_2, \\ 0 & \text{otherwise} \end{cases}$$

Provided that $\mathbf{X}_{t+ni} \in \mathcal{R}_2$, where $\mathcal{R}_2 \subset \mathcal{R}_1$, and assuming a small ϵ_1 and a sufficiently large M , it is not hard to see that if $\mathcal{A}_{t+ni,n} = 1$, then the following hold as a result of (30) and (31). First, $\mathbf{X}_{t+n(i+1)} \in \mathcal{R}_2$. Second, the distance of the vector $\mathbf{X}_{t+n(i+1)}$ from the boundary of \mathcal{R}_2 becomes the distance of \mathbf{X}_{t+ni} plus at least $n\delta_{\mathcal{A}}$. Third,

$$\|\mathbf{X}_{t+n(i+1)}\| \geq \|\mathbf{X}_{t+ni}\| + n\delta_{\mathcal{A}},$$

where $\delta_{\mathcal{A}}$ is an appropriately small positive constant. Fig. 5 shows the region \mathcal{R}_2 rotated clockwise, and provides examples for the case where $\mathcal{A}_{t+ni,n} = 1$. Specifically, when \mathbf{X}_{t+ni} equals one of the points \mathbf{X} , \mathbf{X}_1 , and \mathbf{X}_2 , the figure assumes the

drift vector $\Delta_{t+ni,n}$ is within the ϵ_2 -neighborhood of $\bar{\Delta}_{\mathbf{X}_{t+ni}}$. For points \mathbf{X}_1 and \mathbf{X}_2 , the figure also shows the increases in their distance from the boundary of \mathcal{R}_2 , and denotes them by d_1 and d_2 , respectively. These values, as mentioned above, are lower-bounded by $n\delta_{\mathcal{A}}$ as a result of (30) and (31). To see this note that, as shown in the figure and suggested by (30), when \mathbf{X}_{t+ni} deviation from \mathbf{X}_{min} , i.e., when it deviates from the central line in the figure, the vector $\bar{\Delta}_{\mathbf{X}_{t+ni}}$ gets a component towards the central line. This and the assumption that the angular deviations in the ϵ_2 -neighborhoods are less than half of the one defining region \mathcal{R}_2 , as assumed in (31), ensure that after n steps the backlog vector remains in \mathcal{R}_2 , and that the distance from the boundary of \mathcal{R}_2 increases when $\mathcal{A}_{t+ni,n} = 1$. Using a similar argument, it is easy to see that when ϵ_1 is small, an event of the type $\mathcal{A}_{t+ni,n} = 1$ increases the norm of backlog vector more than $n\delta_{\mathcal{A}}$. On the other hand, if $\mathcal{A}_{t+ni,n} = 0$ with at most probability ζ , both the distance of $\mathbf{X}_{t+n(i+1)}$ from \mathcal{R}_2 and $\|\mathbf{X}_{t+n(i+1)}\|$, compared to the distance of \mathbf{X}_{t+ni} and $\|\mathbf{X}_{t+ni}\|$, respectively, decrease at most by $n\sqrt{2}(A_{max} + D_{max})$. In Fig. 5, the point \mathbf{X}_3 is an example of this case, where the vector $\Delta_{t+ni,n}$ can be anywhere inside the outer circle, centered at \mathbf{X}_3 , but outside the inner circle defining the ϵ_2 -neighborhood of the vector $\mathbf{X}_3 + n\bar{\Delta}_{\mathbf{X}_3}$.

In the rest of the proof, as the worst case, we assume that for $\mathbf{X}_{t+ni} \in \mathcal{R}_2$, $i \geq 0$, the event $\{\mathcal{A}_{t+ni,n} = 1\}$ occurs with probability $(1 - \zeta)$. Note that $\mathcal{R}_2 \subset \mathcal{R}_1$, and when $\mathbf{X}_{t+ni} \in \mathcal{R}_2$, the inequality (29) holds regardless of the past history \mathcal{H}_{t+ni} . Let the event that $\mathcal{A}_{t+ni,n} = 1$ be a success. Based on the previous assumption, for $\mathbf{X}_{t+ni} \in \mathcal{R}_2$, this success event occurs with probability $(1 - \zeta)$ regardless of the past. Now consider the sequence $\{\mathbf{X}_{t+ni}\}$, $0 \leq i \leq m-1$, and let $m_{(1-\zeta)}$ be the number successes of the type $\{\mathcal{A}_{t+ni,n} = 1\}$ out of the m associated trials. The above observations imply that if $\mathbf{X}_{t+ni} \in \mathcal{R}_2$, for $0 \leq i \leq m-1$, and if

$$(m - m_{(1-\zeta)})\sqrt{2}(A_{max} + D_{max}) < m_{(1-\zeta)}\delta_{\mathcal{A}},$$

then $\mathbf{X}_{t+nm} \in \mathcal{R}_2$, and

$$\begin{aligned} \|\mathbf{X}_{t+nm}\| &\geq \|\mathbf{X}_t\| + m_{(1-\zeta)}n\delta_{\mathcal{A}} \\ &\quad - (m - m_{(1-\zeta)})n\sqrt{2}(A_{max} + D_{max}). \end{aligned}$$

Using the above, we see that a sufficient condition for the sequence $\{\mathbf{X}_{t+nm}, m \geq 1\}$ to stay within \mathcal{R}_2 , and

$$\|\mathbf{X}_{t+nm}\| \geq \|\mathbf{X}_t\| + n m \epsilon_3 (\delta_{\mathcal{A}} + \sqrt{2}(A_{max} + D_{max})), \quad (33)$$

for some ϵ_3 with

$$\epsilon_3 < \frac{\delta_{\mathcal{A}}}{\sqrt{2}(A_{max} + D_{max}) + \delta_{\mathcal{A}}}$$

is that for all $m \geq 1$,

$$\begin{aligned} r_m &\triangleq 1 - \frac{m_{(1-\zeta)}}{m} \\ &< \frac{\delta_{\mathcal{A}}}{\sqrt{2}(A_{max} + D_{max}) + \delta_{\mathcal{A}}} - \epsilon_3. \end{aligned} \quad (34)$$

In what follows, we show that with positive probability the above inequality holds for all $m \geq 1$.

Starting at $M\mathbf{X}_{min}$, let $\tau_{\mathcal{R}_2}$ be the first time that the ratio r_m does not satisfy (34). Consider the sequence

$$\{\mathcal{A}_{t+nm,n}, 0 \leq m \leq \tau_{\mathcal{R}_2} - 2\}. \quad (35)$$

For $0 \leq m \leq \tau_{\mathcal{R}_2} - 1$, the discussion leading to (33) and (34) implies that $\mathbf{X}_{t+nm} \in \mathcal{R}_2$. Furthermore, this discussion shows that the sequence can be considered as a truncated Bernoulli process with success probability $(1 - \zeta)$. An intuitive yet important observation is that for an infinite sequence of Bernoulli trials $\{B_i, i \geq 0\}$ with success probability $(1 - \zeta)$, for any given $\epsilon_4 > 0$, with positive probability the ratio of failures *never* reaches $\zeta + \epsilon_4$. This is the key to prove $\tau_{\mathcal{R}_2} = \infty$, or equivalently, (34) holds for all $m \geq 1$, with positive probability. Let the notation r_m be re-used as the failure ratio for the infinite Bernoulli process, i.e.,

$$r_m = 1 - \frac{1}{m} \sum_{i=1}^m B_i.$$

Using large deviation results [34], we have

$$P(r_m - \zeta > \epsilon_4) < \rho^m, \quad (36)$$

where

$$\rho = \inf_{s>0} M_{Z_\zeta}(s) < 1,$$

where $Z_\zeta = 1 - B_1 - \zeta - \epsilon_4$, and $M_{Z_\zeta}(s)$ is the characteristic function of Z_ζ . The above inequality indicates that with probability at least $(1 - \rho^m)$, the ratio of failures after m trials is less than or equal to $\zeta + \epsilon_4$.

To further study r_m , we consider the infinite Bernoulli process in a sequence of stages. In the first and second stages, we consider m Bernoulli trials. However, after the second stage, for the i_{th} stage, we consider the next subsequent $2^{i-2}m$ trials. Since trials are independent, with probability $(1 - \zeta)^m$, we can have only successes for the first m trials, and thus, the ratio r_m never goes beyond zero, i.e.,

$$r_j = 0, 1 \leq j \leq m.$$

For the second stage with the next m trials, using (36), we see that with probability at least $(1 - \zeta)^m(1 - \rho^m)$

$$\max_{0 \leq j \leq 2^m} r_j \leq \frac{0 + m(\zeta + \epsilon_4)}{m + m(\zeta + \epsilon_4)} < 2(\zeta + \epsilon_4),$$

where the first inequality refers to the worst case where in the second stage of m trials, the failures happen in the beginning of the stage, i.e., when the $(m+1)_{th}$, $(m+2)_{th}$, ..., and $(m+m(\zeta + \epsilon_3))_{th}$ trials are all failures. Inductively, considering the $(l+2)_{th}$ stage, we see that with probability at least $(1 - \zeta)^m \prod_{p=0}^l (1 - \rho^{2^p m})$

$$\max_{1 \leq j \leq 2^{(l+1)m}} r_j \leq \frac{(2^{l+1} - 1)m(\zeta + \epsilon_4)}{2^l m + 2^l m(\zeta + \epsilon_4)} < 2(\zeta + \epsilon_4), \quad (37)$$

where the numerator is the total number of failures up to the end of $(l+2)_{th}$ stage, and the denominator corresponds to the worst case where the failures in the $(l+2)_{th}$ stage all occur in the beginning of the stage. Therefore, with probability at least

$$p_\zeta = (1 - \zeta)^m \prod_{p=0}^\infty (1 - \rho^{2^p m})$$

the ratio r_m , $m \geq 1$, always stays below $2(\zeta + \epsilon_4)$. But

$$p_\zeta > (1 - \zeta)^m \Pi_{p=0}^\infty (1 - (\rho^m)^{p+1}).$$

This and Lemma 5 indicate that

$$p_\zeta > 0.$$

The above discussion implies that with a positive probability, not less than p_ζ , the ratio r_m associated with the sequence in (35) stays below $2(\zeta + \epsilon_4)$. Hence, if ζ and ϵ_4 are chosen such that

$$\zeta + \epsilon_4 < \frac{1}{2} \left(\frac{\delta_A}{\sqrt{2}(A_{max} + D_{max}) + \delta_A} - \epsilon_3 \right), \quad (38)$$

then starting at $\mathbf{X}_t = M\mathbf{X}_{min}$, the inequality in (34) holds for all $m \geq 1$ with positive probability. Since this latter statement can be generalized to the case where $\mathbf{X}_t \in \mathcal{R}_2$, we have that

$$P(\forall m \geq 0, \mathbf{X}_{t+nm} \in \mathcal{R}_2 \text{ and (33) holds } |\mathbf{X}_t \in \mathcal{R}_2|) > 0 \quad (39)$$

if (38) holds. But (38) can be satisfied since the choice for a positive ϵ_4 is arbitrary, and as mentioned in the discussion leading to (29), ζ can be chosen arbitrarily small. Hence, for an appropriate choice of parameters, (39) holds, which suggests that with positive probability \mathbf{X}_{t+nm} stays in \mathcal{R}_2 , and its norm increases (at least) linearly with m . Since by C4 with positive probability $\mathbf{X}_t \in \mathcal{R}_2$ for some t , and arrivals and departures are bounded implying for $0 \leq j \leq n$, $\|\mathbf{X}_{t+mn-j}\| \geq \|\mathbf{X}_{t+nm}\| - C$, for some $C > 0$, we see that (39) indicates that with positive probability

$$\lim_{i \rightarrow \infty} \|\mathbf{X}_{t+i}\| = \infty.$$

This shows that when the input rate is given by (28), and thus, when it is outside the region $\theta_\infty \Gamma$, for any bounded region \mathcal{C} , with positive probability the process \mathbf{X}_{t+i} never returns to \mathcal{C} , and hence, $\mathbb{E}[\tau_{\mathcal{C}}] = \infty$, implying network instability. This completes the proof of part (c) of the theorem. \blacksquare

APPENDIX

LEMMAS

Lemma 1: For any $\epsilon > 0$, regardless of the past history \mathcal{H}_t , there exists a sufficiently large K_ϵ such that for all $\mathbf{s} \in \mathcal{S}$ and $K > K_\epsilon$

$$\left| \mathbb{E}[(\tau_K + 1)\pi(\mathbf{s}) - \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} | \mathcal{H}_t] \right| < \epsilon \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t].$$

Proof: Since $\tau_{i+1} - \tau_i > 2N_c$, it is easy to verify that $\tau_K \rightarrow \infty$, a.s., as $K \rightarrow \infty$. This almost surely convergence and the ergodicity of channel process, as stated in Section III-B, imply that as $K \rightarrow \infty$

$$\frac{1}{\tau_K + 1} \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} \rightarrow \pi(\mathbf{s}), \text{ a.s.} \quad (40)$$

Moreover, since the channel convergence in Section III-B is uniform in the past history and t , and since the number of channel states is finite, we see that the above convergence is uniform in t , \mathcal{H}_t , and \mathbf{s} . Thus, for any $\epsilon' > 0$ and $\zeta > 0$,

we can find a sufficiently large $K_{\epsilon', \zeta}$ independent of the past history \mathcal{H}_t and \mathbf{s} such that [29]

$$P\left(\sup_{K > K_{\epsilon', \zeta}} \left| \pi(s) - \frac{1}{\tau_K + 1} \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} \right| > \epsilon' | \mathcal{H}_t\right) < \zeta. \quad (41)$$

Given \mathcal{H}_t , let $\mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}$ denote the set of all $\omega \in \Omega$ with the property that

$$\sup_{K > K_{\epsilon', \zeta}} \left| \pi(s) - \frac{1}{\tau_K + 1} \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} \right| > \epsilon'.$$

By (41), we have that

$$P(\mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'} | \mathcal{H}_t) < \zeta. \quad (42)$$

Suppose $K > K_{\epsilon', \zeta}$ and let

$$\Delta = \mathbb{E}\left[(\tau_K + 1)\pi(\mathbf{s}) - \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} | \mathcal{H}_t\right].$$

Using conditional expectations and the definition of $\mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}$, and considering the fact that $0 \leq \pi(\mathbf{s}) \leq 1$ and $\tau_K \geq 0$, we can show that

$$\begin{aligned} \Delta &\leq P(\omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}) \mathbb{E}[(\tau_K + 1) | \omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}, \mathcal{H}_t] \\ &\quad + P(\omega \in \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}) \mathbb{E}[(\tau_K + 1) | \omega \in \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}, \mathcal{H}_t] \end{aligned} \quad (43)$$

Similarly, we obtain

$$\begin{aligned} &\mathbb{E}[(\tau_K + 1) | \mathcal{H}_t] \\ &= P(\omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}) \mathbb{E}[(\tau_K + 1) | \omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}, \mathcal{H}_t] \\ &\quad + P(\omega \in \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}) \mathbb{E}[(\tau_K + 1) | \omega \in \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}, \mathcal{H}_t]. \end{aligned}$$

Since $\tau_K \geq 0$, the above implies that

$$\begin{aligned} &P(\omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}) \mathbb{E}[(\tau_K + 1) | \omega \notin \mathcal{A}_{K_{\epsilon', \zeta}, \epsilon'}, \mathcal{H}_t] \\ &\leq \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t] \end{aligned} \quad (44)$$

In addition, w.p.1, $\tau_K + 1 \leq (K+1)(1+L_1)N_c$. It thus follows from (42), (43), and (44) that

$$\Delta \leq \epsilon' \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t] + \zeta(K+1)(1+L_1)N_c$$

Noting the fact that $\tau_K \geq 2KN_c$, we obtain

$$\begin{aligned} \Delta &\leq \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t] \left(\epsilon' + \zeta \frac{(K+1)(1+L_1)N_c}{\mathbb{E}[(\tau_K + 1) | \mathcal{H}_t]} \right) \\ &< \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t] \left(\epsilon' + \zeta \frac{(K+1)(1+L_1)N_c}{2KN_c + 1} \right) \\ &= \epsilon \mathbb{E}[(\tau_K + 1) | \mathcal{H}_t], \end{aligned}$$

where

$$\epsilon = \epsilon' + \zeta \frac{(K+1)(1+L_1)N_c}{2KN_c + 1}$$

can be made arbitrarily small by choosing sufficiently small values for ϵ' and ζ . A similar discussion holds for $-\Delta$ with the same ϵ , completing the proof. \blacksquare

Lemma 2: For any given $\epsilon > 0$, there exists a sufficiently large constant $K_\epsilon > 0$ such that for all $K > K_\epsilon$, we can find a proper $M_{\epsilon, K}$ such that if $\|\mathbf{X}_t\| > M_{\epsilon, K}$, the following holds

$$\begin{aligned} & \mathbb{E}\left[\sum_{k=0}^{\tau_K} \mathbf{X}_t \mathbf{a} - \mathbf{X}_{t+k} \mathbf{D}_{t+k}^* \mid \mathcal{H}_t\right] \\ & \leq \mathbb{E}\left[(\tau_K + 1)\|\mathbf{X}_t\| \left(\epsilon - \left(1 - \max_{\mathbf{s} \in \mathcal{S}} \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}}\right) \chi(\mathbf{X}_t)\right) \mid \mathcal{H}_t\right] \end{aligned}$$

Proof: To prove the lemma, we first note that by the definition of $\mathbf{D}^*(\mathbf{X}_t, \mathbf{s})$, and the assumption that departures are bounded by D_{max} , we have

$$\begin{aligned} \mathbf{X}_{t+k} \mathbf{D}_{t+k}^* &= \max_{\mathbf{I} \in \mathcal{I}} \mathbf{X}_{t+k} \mathbf{D}(\mathbf{s}_{t+k}, \mathbf{I}) \\ &\geq \max_{\mathbf{I} \in \mathcal{I}} \mathbf{X}_t \mathbf{D}(\mathbf{s}_{t+k}, \mathbf{I}) - \max_{\mathbf{I} \in \mathcal{I}} \left(\sum_{i=0}^{k-1} \mathbf{D}_{t+i} \mathbf{D}(\mathbf{s}_{t+k}, \mathbf{I})\right) \\ &\geq \max_{\mathbf{I} \in \mathcal{I}} \mathbf{X}_t \mathbf{D}(\mathbf{s}_{t+k}, \mathbf{I}) - kND_{max}^2. \end{aligned} \quad (45)$$

Using (11), we also observe that

$$\begin{aligned} & \mathbf{X}_t \mathbf{a} - \mathbb{E}_{\mathbf{s}}[\mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s})] \\ &= \mathbf{X}_t \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}} \mathbf{D}(\mathbf{s}, \mathbf{I}) - \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \\ &= \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \left(\sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}} \left(\mathbf{X}_t \mathbf{D}(\mathbf{s}, \mathbf{I}) - \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \right) \right. \\ &\quad \left. - \left((1 - \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}}) \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \right) \right). \end{aligned} \quad (46)$$

Since by definition for all $\mathbf{I} \in \mathcal{I}$

$$\mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \geq \mathbf{X}_t \mathbf{D}(\mathbf{s}, \mathbf{I}),$$

we have

$$\begin{aligned} & \mathbf{X}_t \mathbf{a} - \mathbb{E}_{\mathbf{s}}[\mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s})] \\ &\leq - \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \left(1 - \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}}\right) \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \\ &\leq - \left(1 - \max_{\mathbf{s} \in \mathcal{S}} \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}}\right) \sum_{\mathbf{s} \in \mathcal{S}} \pi(\mathbf{s}) \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \\ &= -\|\mathbf{X}_t\| \left(1 - \max_{\mathbf{s} \in \mathcal{S}} \sum_{\mathbf{I} \in \mathcal{I}} \beta_{\mathbf{s}, \mathbf{I}}\right) \chi(\mathbf{X}_t), \end{aligned} \quad (47)$$

where the last equality follows from the definition of $\chi(\mathbf{X}_t)$.

Back to the inequality in the lemma, using (45), we have

$$\begin{aligned} & \mathbb{E}\left[\sum_{k=0}^{\tau_K} \mathbf{X}_t \mathbf{a} - \mathbf{X}_{t+k} \mathbf{D}_{t+k}^* \mid \mathcal{H}_t\right] \leq \mathbb{E}\left[(\tau_K + 1)^2 ND_{max}^2 \mid \mathcal{H}_t\right] \\ & \quad + \mathbb{E}\left[\sum_{k=0}^{\tau_K} \left(\mathbf{X}_t \mathbf{a} - \mathbf{X}_t \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} \mathbf{D}^*(\mathbf{X}_t, \mathbf{s})\right) \mid \mathcal{H}_t\right] \\ &= \mathbb{E}\left[(\tau_K + 1)^2 ND_{max}^2 \mid \mathcal{H}_t\right] \\ & \quad + \mathbb{E}\left[\sum_{k=0}^{\tau_K} \mathbf{X}_t \mathbf{a} - \mathbf{X}_t \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \sum_{k=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+k)=\mathbf{s}} \mid \mathcal{H}_t\right]. \end{aligned}$$

Using Lemma 1, for $\epsilon_1 > 0$ and sufficiently large K_1 , we have that for $K > K_1$

$$\begin{aligned} & \mathbb{E}\left[\sum_{k=0}^{\tau_K} \mathbf{X}_t \mathbf{a} - \mathbf{X}_{t+k} \mathbf{D}_{t+k}^* \mid \mathcal{H}_t\right] \\ & \leq \mathbb{E}\left[(\tau_K + 1)^2 ND_{max}^2 \mid \mathcal{H}_t\right] + \mathbb{E}\left[(\tau_K + 1) \mathbf{X}_t \mathbf{a} \right. \\ & \quad \left. - (\tau_K + 1) \mathbf{X}_t \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) (\pi(\mathbf{s}) - \epsilon_1) \mid \mathcal{H}_t\right] \\ &= \mathbb{E}\left[(\tau_K + 1)^2 ND_{max}^2 \mid \mathcal{H}_t\right] \\ & \quad + \epsilon_1 \mathbb{E}\left[(\tau_K + 1) \|\mathbf{X}_t\| |\mathcal{S}| \sqrt{ND_{max}} \mid \mathcal{H}_t\right] \\ & \quad + \mathbb{E}\left[(\tau_K + 1) \left(\mathbf{X}_t \mathbf{a} - \mathbb{E}_{\mathbf{s}}[\mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s})]\right) \mid \mathcal{H}_t\right]. \end{aligned} \quad (48)$$

Combining (47) and (48), we obtain the inequality in lemma with

$$\epsilon = \frac{(\tau_K + 1) ND_{max}^2}{\|\mathbf{X}_t\|} + \epsilon_1 |\mathcal{S}| \sqrt{ND_{max}}.$$

The choice for a positive ϵ is arbitrary since one can first select $K_\epsilon \geq K_1$ such that for all $K > K_\epsilon$, ϵ_1 is sufficiently small. After selecting K , because w.p.1 $\tau_K + 1 \leq (K+1)(1+L_1)N_c$, one can choose $M_{\epsilon, K}$ such that for $\|\mathbf{X}_t\| > M_{\epsilon, K}$ the first term in ϵ is also sufficiently small, completing the proof. \blacksquare

Lemma 3: For any given $\epsilon > 0$, there exists a sufficiently large constant $K_\epsilon > 0$ such that for all $K > K_\epsilon$, we can find a proper $M_{\epsilon, K}$ such that if $\|\mathbf{X}_t\| > M_{\epsilon, K}$, the following holds

$$\begin{aligned} & \mathbb{E}\left[\sum_{i=0}^{\tau_K} \mathbf{X}_{t+i} \mathbf{D}_{t+i}^* \mid \mathcal{H}_t\right] \\ & \leq \mathbb{E}\left[(\tau_K + 1) \|\mathbf{X}_t\| (\chi(\mathbf{X}_t) + \epsilon) \mid \mathcal{H}_t\right]. \end{aligned}$$

Proof: Using the definition of $\mathbf{D}^*(\mathbf{X}, \mathbf{s})$, for the LHS of the inequality in the lemma we can show that

$$\text{LHS} = \mathbb{E}\left[\sum_{i=0}^{\tau_K} \max_{\mathbf{I} \in \mathcal{I}} \left(\left(\mathbf{X}_t + \sum_{j=0}^{i-1} (\mathbf{A}_{t+j} - \mathbf{D}_{t+j} + \mathbf{U}_{t+j}) \right) \mathbf{D}(\mathbf{s}_{t+i}, \mathbf{I}) \right) \mid \mathcal{H}_t\right].$$

Since arrivals and departures are bounded by A_{max} and D_{max} , respectively, we have that

$$\begin{aligned} \text{LHS} &\leq \mathbb{E}\left[\sum_{i=0}^{\tau_K} \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}_{t+i}) \mid \mathcal{H}_t\right] \\ &+ \mathbb{E}\left[\sum_{i=0}^{\tau_K} i N A_{max} D_{max} + \sum_{i=0}^{\tau_K} i N D_{max}^2 \mid \mathcal{H}_t\right]. \end{aligned} \quad (49)$$

Let Σ be the first term of the RHS of the above inequality. We have

$$\begin{aligned} \Sigma &= \mathbb{E}\left[\sum_{i=0}^{\tau_K} \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \mathbf{1}_{\mathbf{s}(t+i)=\mathbf{s}} \mid \mathcal{H}_t\right] \\ &= \mathbf{X}_t \sum_{\mathbf{s} \in \mathcal{S}} \mathbf{D}^*(\mathbf{X}_t, \mathbf{s}) \mathbb{E}\left[\sum_{i=0}^{\tau_K} \mathbf{1}_{\mathbf{s}(t+i)=\mathbf{s}} \mid \mathcal{H}_t\right]. \end{aligned}$$

Using Lemma 1, for any positive ϵ_1 , we can find a sufficiently large K_1 such that for $K > K_1$

$$\begin{aligned}\Sigma &\leq \mathbf{X}_t \sum_{s \in \mathcal{S}} \mathbf{D}^*(\mathbf{X}_t, s) \mathbb{E} \left[(\tau_K + 1)(\pi(s) + \epsilon_1) \mid \mathcal{H}_t \right] \\ &= \mathbb{E} \left[(\tau_K + 1) \sum_{s \in \mathcal{S}} \pi(s) \mathbf{X}_t \mathbf{D}^*(\mathbf{X}_t, s) \right. \\ &\quad \left. + \epsilon_1 |\mathcal{S}| \sqrt{N} \|\mathbf{X}_t\| D_{max} (\tau_K + 1) \mid \mathcal{H}_t \right] \\ &= \mathbb{E} \left[(\tau_K + 1) \|\mathbf{X}_t\| \left(\chi(\mathbf{X}_t) + \epsilon_1 \sqrt{N} |\mathcal{S}| D_{max} \right) \mid \mathcal{H}_t \right],\end{aligned}\quad (50)$$

where the last equality follows from the definition of $\chi(\mathbf{X}_t)$.

Considering inequalities (49) and (50), we obtain

$$\text{LHS} \leq \mathbb{E} \left[(\tau_K + 1) \|\mathbf{X}_t\| (\chi(\mathbf{X}_t) + \epsilon_2) \mid \mathcal{H}_t \right], \quad (51)$$

where

$$\epsilon_2 = \epsilon_1 |\mathcal{S}| \sqrt{N} D_{max} + (\tau_K + 1) N \frac{A_{max} D_{max} + D_{max}^2}{\|\mathbf{X}_t\|}.$$

To complete the proof, it remains to show that ϵ_2 can be made arbitrarily small. Consider any positive ϵ . We first choose K_ϵ such that for $K > K_\epsilon$ the value of ϵ_1 is sufficiently small to make the first term in ϵ_2 less than $\frac{\epsilon}{2}$. Since $\tau_k + 1 \leq (K+1)(1+L_1)N_c$, we see that for a given K with $K > K_\epsilon$ if $\|\mathbf{X}_t\| > M_{\epsilon,K}$ for a sufficiently large $M_{\epsilon,K}$, then the second term in ϵ_2 can also be less than $\frac{\epsilon}{2}$. Therefore, for any positive ϵ , if $K > K_\epsilon$ and $\|\mathbf{X}_t\| > M_{\epsilon,K}$, for appropriate values of K_ϵ and $M_{\epsilon,K}$, then the inequality (51) holds with $\epsilon_2 < \epsilon$. But this means the inequality also holds for ϵ , as required. ■

Lemma 4: Suppose $6\theta_\varphi < \alpha$, and let ϵ be a positive real number. For any given ϵ , there exists a constant K_ϵ such that if $K > K_\epsilon$, then for $\|\mathbf{X}_t\| > M_{\epsilon,K}$ the following holds

$$\begin{aligned}\mathbb{E} \left[\sum_{i=\tau_0+1}^{\tau_K} \mathbf{X}_{t+i} \mathbf{D}_{t+i} \mid \mathcal{H}_t \right] \\ \geq \mathbb{E} \left[(\tau_K + 1) \|\mathbf{X}_t\| (R_\infty (\tilde{\phi}(t) - \alpha - 3\theta_\varphi) - \epsilon) \mid \mathcal{H}_t \right],\end{aligned}$$

where $M_{\epsilon,K}$ is a sufficiently large constant depending on ϵ and K , and R_∞ is defined in (8).

Proof: The essence of the proof in this lemma is finding a lower-bound for the percentage of time that *near optimal* values for N_1 are used by DCP. We prove that this percentage is close to R_∞ . First, we place a requirement on $\|\mathbf{X}_t\|$ for a given K . Later in the proof, we find an appropriate lower-bound K_ϵ for K according to the value of ϵ . Note that w.p.1, for any given K , $\tau_K \leq (K+1)(1+L_1)N_c$. Therefore, since departures and arrivals are bounded by D_{max} and A_{max} , respectively, we can easily see that for $0 \leq i \leq \tau_K$, $\|\mathbf{X}_{t+i} - \mathbf{X}_t\| < C'_K$, where C'_K is an appropriate constant depending on K . Having this inequality, we can find an appropriate constant M'_K , depending on K , such that if

$$\|\mathbf{X}_t\| > M'_K, \quad (52)$$

then the following statements hold according to Property 1 and Property 2, respectively, with $\epsilon_1 < \frac{1}{2}(\frac{\alpha}{6} - \theta_\varphi)$.

Statement 1: For $t \leq t_1 \leq t + \tau_K$, $t \leq t_2 \leq t + \tau_K$, and any $N_1 \in \mathcal{N}_1$,

$$|\phi(\mathbf{X}_{t_1}, N_1) - \phi(\mathbf{X}_{t_2}, N_1)| < \epsilon_1. \quad (53)$$

Statement 2: For any τ_i , with $0 \leq i \leq K$, and any $N_1 \in \mathcal{N}_1$, with probability $(1 - \varrho_\varphi)$, and regardless of i and the past history at time $t + \tau_i + 1$, $\mathcal{H}_{t+\tau_i+1}$, we have

$$\begin{aligned}|\varphi^r(t + \tau_i + 1) - \\ \phi(t + \tau_i + 1, N_1^r(t + \tau_i + 1))| < \theta_\varphi + \epsilon_1.\end{aligned}\quad (54)$$

Similarly, with probability $(1 - \varrho_\varphi)$, and regardless of i and the past history at time $t + \tau_i + 1 + N_c$, $\mathcal{H}_{t+\tau_i+1+N_c}$, we have

$$\begin{aligned}|\varphi(t + \tau_i + 1) - \\ \phi(t + \tau_i + 1 + N_c, N_1(t + \tau_i + 1))| < \theta_\varphi + \epsilon_1.\end{aligned}\quad (55)$$

Remark 1: Property 2 states inequalities in Statement 2 may hold in general with *different* probabilities all not less than $(1 - \varrho_\varphi)$. However, to consider the *worst case* analysis, in Statement 2, we have assumed these inequalities, with the given conditions, hold with the *same* probability $(1 - \varrho_\varphi)$ for all i , where $0 \leq i \leq K$.

Remark 2: Consider the $i+1$ th and the $j+1$ th rounds, where $0 \leq i, j \leq K$ and $i \neq j$. Since inequalities (54) and (55) in Statement 2 may hold in the $i+1$ th round with probability $1 - \varrho_\varphi$ regardless of $\mathcal{H}_{t+\tau_i+1}$ and $\mathcal{H}_{t+\tau_i+1+N_c}$, respectively, Statement 2 implies that the event that (54) or the one that (55) holds in the $i+1$ th round is independent of the inequality (54) or (55) holding in the $j+1$ th round. In addition, the event that (54) holds in the $i+1$ th round is independent of (55) holding in the same round.

Before going to the main part of the proof, we first derive two key inequalities. To obtain the first one, note that for any two time instants t_1 and t_2 , with $t \leq t_1 \leq t + \tau_K$ and $t \leq t_2 \leq t + \tau_K$, using (53), we have that

$$|\phi(t_1, \tilde{N}_1(\mathbf{X}_{t_1})) - \phi(t_2, \tilde{N}_1(\mathbf{X}_{t_1}))| < \epsilon_1, \quad (56)$$

and

$$|\phi(t_1, \tilde{N}_1(\mathbf{X}_{t_2})) - \phi(t_2, \tilde{N}_1(\mathbf{X}_{t_2}))| < \epsilon_1.$$

By the definition of $\tilde{N}_1(\mathbf{X})$ and the inequality in (56), we have

$$\begin{aligned}\phi(t_1, \tilde{N}_1(\mathbf{X}_{t_1})) - \phi(t_2, \tilde{N}_1(\mathbf{X}_{t_2})) \\ \leq \phi(t_1, \tilde{N}_1(\mathbf{X}_{t_1})) - \phi(t_2, \tilde{N}_1(\mathbf{X}_{t_1})) < \epsilon_1\end{aligned}$$

We can obtain the other direction of the inequality similarly. Thus,

$$|\phi(t_1, \tilde{N}_1(\mathbf{X}_{t_1})) - \phi(t_2, \tilde{N}_1(\mathbf{X}_{t_2}))| < \epsilon_1. \quad (57)$$

This inequality shows that when backlog vector has a large absolute value, the optimal ϕ does not vary significantly in a limited time horizon. In particular, the variation approaches zero when $\|\mathbf{X}_t\|$ approaches ∞ .

To derive the second key inequality, first note that based on the definition of τ_i given in the proof of part(a) of the theorem, the $i+1$ th round after time t begins at $t + \tau_i + 1$, and the time interval between $t + \tau_0 + 1$ and $t + \tau_K + 1$ consists of K scheduling rounds. To simplify the notation, let \tilde{N}_1 be

the optimal value of N_1 for the first round after time t , i.e., $\tilde{N}_1 = \tilde{N}_1(\mathbf{X}_{t+\tau_0+1})$. In addition, let $N_1^r(j)$ be the candidate value for N_1 in the $j+1$ th round, and let $N_1(j)$ be the value of N_1 used in the update interval of the $j+1$ th round, i.e., $N_1^r(j) = N_1^r(t + \tau_j + 1)$, and $N_1(j) = N_1(t + \tau_j + 1)$.

Now, consider the $i+1$ th round, $i \geq 0$, and suppose the optimal N_1 is selected at this round, i.e., $N_1^r(i) = \tilde{N}_1(\mathbf{X}_{t+\tau_i+1})$. Let $\hat{N}_1 = \hat{N}_1(\mathbf{X}_{t+\tau_i+1})$. Then the inequality in (54) and the preceding inequality imply that with probability $(1 - \varrho_\varphi)$

$$|\varphi^r(t + \tau_i + 1) - \phi(t + \tau_0 + 1, \tilde{N}_1)| < 2\epsilon_1 + \theta_\varphi. \quad (58)$$

Let

$$\epsilon' = 2\epsilon_1 + \theta_\varphi. \quad (59)$$

Based on the assumption $6\theta_\varphi < \alpha$ imposed by the Lemma and that $\epsilon_1 < \frac{1}{2}(\frac{\alpha}{6} - \theta_\varphi)$, we have

$$0 < 6\epsilon' < 6\left(\left(\frac{\alpha}{6} - \theta_\varphi\right) + \theta_\varphi\right) = \alpha. \quad (60)$$

The inequality (58) is the second key inequality required for the rest of the proof.

We are now in a position to explain the essence of the proof, where we find a lower-bound for the fraction of time in the horizon of K rounds in which near optimal values for N_1 are used. Towards this end, we first assume that the inequalities in (54) and (55) hold with probability one for all K scheduling rounds, thus assuming $\varrho_\varphi = 0$ in Statement 2. We then extend our discussion to realistic cases where $\varrho_\varphi > 0$.

Discussion assuming $\varrho_\varphi = 0$: Suppose at the $i+1$ th round, $i \geq 1$, the optimal N_1 corresponding to $\mathbf{X}_{t+\tau_i+1}$ is selected, i.e., $N_1^r(i) = \hat{N}_1 = \hat{N}_1(\mathbf{X}_{t+\tau_i+1})$. Considering the scheduling policy, with respect to the update of N_1 in $i+1$ th scheduling round, there are two possible cases:

Case 1: In this case, we assume $\varphi^r(t + \tau_i + 1) > \varphi(t + \tau_{i-1} + 1) + \alpha$. Thus, according to the update rule, N_1 gets updated at the $i+1$ th round, and takes the value $N_1(i) = N_1^r(i) = \hat{N}_1$. However, it remains unchanged until the $K+1$ th round. We can prove this statement by induction. To see this, assume that N_1 remains fixed after the $i+1$ th but changes for the first time in the j th + 1 round, where $j > i$. Therefore, by the update rule, we must have

$$\varphi^r(t + \tau_j + 1) > \varphi(t + \tau_{j-1} + 1) + \alpha. \quad (61)$$

Since

$$|\phi(t + \tau_j + 1, N_1^r(j)) - \phi(t + \tau_0 + 1, N_1^r(j))| < \epsilon_1,$$

and

$$|\varphi^r(t + \tau_j + 1) - \phi(t + \tau_j + 1, N_1^r(j))| < \theta_\varphi + \epsilon_1,$$

which follow from (53) and (54), respectively, and the assumption that $\varrho_\varphi = 0$, we have

$$\begin{aligned} \varphi^r(t + \tau_j + 1) &< \phi(t + \tau_0 + 1, N_1^r(j)) + 2\epsilon_1 + \theta_\varphi \\ &\leq \phi(t + \tau_0 + 1, \tilde{N}_1) + \epsilon', \end{aligned} \quad (62)$$

where the last inequality follows from the definition of \tilde{N}_1 .

Similarly, since by assumption $N(j-1) = \hat{N}_1$, we can use (53) and (55) to show that

$$|\varphi(t + \tau_{j-1} + 1) - \phi(t + \tau_i + 1, \hat{N}_1)| < \epsilon'.$$

Considering this inequality and (57), we obtain

$$\varphi(t + \tau_{j-1} + 1) > \phi(t + \tau_0 + 1, \tilde{N}_1) - \epsilon' - \epsilon_1. \quad (63)$$

Finally, considering (61), (62), and (63), we obtain

$$\begin{aligned} \phi(t + \tau_0 + 1, \tilde{N}_1) + \epsilon' &> \\ \phi(t + \tau_0 + 1, \tilde{N}_1) - \epsilon' - \epsilon_1 + \alpha, \end{aligned}$$

which implies that $2\epsilon' + \epsilon_1 > \alpha$. This is in contradiction with (60) stating that $6\epsilon' < \alpha$. Therefore, $N_1(j) = \hat{N}_1$ for $i \leq j \leq K-1$, proving the claim.

A byproduct of the above discussion is that after the i th round, $\varphi(t + \tau_j + 1)$ stays close to $\phi(t + \tau_0 + 1, \tilde{N}_1)$. More precisely, since $N_1(j) = \hat{N}_1$ for $i \leq j \leq K-1$, we have

$$|\varphi(t + \tau_j + 1) - \phi(t + \tau_j + 1, \hat{N}_1))| < \theta_\varphi + \epsilon_1.$$

Moreover, from (53) we have

$$|\phi(t + \tau_j + 1, \hat{N}_1) - \phi(t + \tau_i + 1, \hat{N}_1)| < \epsilon_1.$$

Using the last two inequalities and (57), for $i \leq j \leq K-1$, we obtain

$$|\varphi(t + \tau_j + 1) - \phi(t + \tau_0 + 1, \tilde{N}_1)| < \epsilon' + \epsilon_1, \quad (64)$$

which shows how close is $\varphi(t + \tau_j + 1)$ to $\phi(t + \tau_0 + 1, \tilde{N}_1)$.

Case 2: In this case, we assume $\varphi^r(t + \tau_i + 1) \leq \varphi(t + \tau_{i-1} + 1) + \alpha$. Taking similar steps as in Case 1, we can show that

$$\phi(t + \tau_0 + 1, \tilde{N}_1) - \epsilon' \leq \varphi^r(t + \tau_i + 1)$$

and

$$\varphi(t + \tau_{i-1} + 1) \leq \phi(t + \tau_0 + 1, N_1(i-1)) + \epsilon',$$

Hence, using the assumption, we obtain

$$\phi(t + \tau_0 + 1, \tilde{N}_1) - 2\epsilon' - \alpha \leq \phi(t + \tau_0 + 1, N_1(i-1)). \quad (65)$$

We next show that N_1 gets updated at most once in the rest of $K-(i+1)$ rounds. Let the j_1+1 th round, for $i < j_1 \leq K-1$, be the first round after the $i+1$ th round that N_1 gets updated. Using similar arguments as the ones in Case 1, we have

$$\varphi^r(t + \tau_{j_1} + 1) < \phi(t + \tau_0 + 1, N_1^r(j_1)) + \epsilon'.$$

and

$$\varphi(t + \tau_{j_1-1} + 1) > \phi(t + \tau_0 + 1, N_1(i-1)) - \epsilon',$$

where in the above we have used the assumption that N_1 does not change before the j_1+1 round, and thus have set $N_1(j_1-1) = N_1(i-1)$. Since N_1 gets updated at the j_1+1 th, we have $N_1(j) = N_1^r(j)$. Using this, the update rule, and the last two inequalities, we have

$$\begin{aligned} \phi(t + \tau_0 + 1, N_1(i-1)) - \epsilon' + \alpha & \\ &< \phi(t + \tau_0 + 1, N_1(j_1)) + \epsilon' \end{aligned}$$

This inequality and (65) yield

$$\phi(t + \tau_0 + 1, \tilde{N}_1) - 4\epsilon' < \phi(t + \tau_0 + 1, N_1(j_1)). \quad (66)$$

Similarly, if there exists $j_1 < j_2 \leq K-1$, such that at $t + \tau_{j_2} + 1$, N_1 becomes updated for the second time, we can show that

$$\begin{aligned} \phi(t + \tau_0 + 1, N_1(j_1)) - \epsilon' + \alpha \\ < \phi(t + \tau_0 + 1, N_1(j_2)) + \epsilon'. \end{aligned}$$

In other words,

$$\begin{aligned} \phi(t + \tau_0 + 1, N_1(j_1)) + \alpha - 2\epsilon' \\ < \phi(t + \tau_0 + 1, N_1(j_2)). \end{aligned}$$

Therefore, every time that N_1 becomes updated, the algorithm finds a better estimate for $\phi(t + \tau_0 + 1, \tilde{N}_1)$. More specifically, after each update, the gap between $\phi(t + \tau_0 + 1, N_1(j_k))$ and $\phi(t + \tau_0 + 1, \tilde{N}_1)$ is decreased by $(\alpha - 2\epsilon') > \frac{2}{3}\alpha$. However, (65) shows that the initial gap is $\alpha + 2\epsilon'$, which is less than or equal to $\frac{4}{3}\alpha$. Therefore, N_1 can be updated at most once in the rest of $K - i - 1$ scheduling rounds.

In this case, similar to what we observed in Case 1, $\varphi(t + \tau_j + 1)$ stays close to $\phi(t + \tau_0 + 1, \tilde{N}_1)$. To see this, consider a scheduling round, e.g., $j + 1$ round for $i \leq j \leq K-1$ round, where $N_1(i-1)$ is used. By (53) and (55), we have

$$|\varphi(t + \tau_j + 1) - \phi(t + \tau_0 + 1, N_1(i-1))| < \epsilon'.$$

Considering the above inequality and (65), we obtain

$$|\varphi(t + \tau_j + 1) - \phi(t + \tau_0 + 1, \tilde{N}_1)| < \alpha + 3\epsilon'. \quad (67)$$

In the same manner, if instead of $N_1(i-1)$ an updated version of N_1 is used in an scheduling round, we can use the inequality in (66) to show that the above inequality still holds. Hence, the inequality in (67) holds for all j with $i \leq j \leq K-1$ since, as proved earlier, N_1 becomes updated at most once.

Combining the inequality (64) associated with Case 1 and the inequality (67) associated with Case 2, we see that regardless of which case happens, the following holds for $i \leq j \leq K-1$:

$$|\varphi(t + \tau_j + 1) - \phi(t + \tau_0 + 1, \tilde{N}_1)| < \gamma, \quad (68)$$

where

$$\gamma = \alpha + 3\epsilon'. \quad (69)$$

Inspired by the above inequality, we now define a new random variable R_K as the percentage of time that “near optimal” solution is used in the time horizon consisting of K rounds. By near optimal in a scheduling round, e.g., the $j + 1$ round, we mean a choice of N_1 that ensures $\varphi(t + \tau_j + 1)$ is close to $\phi(t + \tau_0 + 1, \tilde{N}_1)$ in the sense of (68). Intuitively, a larger R_K results in a larger scaling factor, and thus, a better throughput performance. In the following, using the preceding discussions provided in Case 1 and Case 2, we find a lower bound for R_K .

As explained in Section IV-B, in the beginning of each round, e.g., the $j + 1$ round, the optimal N_1 , corresponding to $\mathbf{X}_{t+\tau_j+1}$, is chosen independently with probability δ . Therefore, we see that with probability $(1 - \delta)^{i-1}\delta$, after the first round, the optimal solution is selected for the first time in

the $i + 1$ round, $i \geq 1$. Suppose this event happens at $i + 1$ round, $i \geq 1$. If Case 1 happens, we can partition the time interval between $t + \tau_0 + 1$ and $t + \tau_K + 1$ into three sets. The first set consists of all test intervals. The second set consists of the update intervals before the $i + 1$ round. Finally, the third set consists of the update intervals after the i round. Considering these sets in sequence, we can express the total number of timeslots between $t + \tau_0 + 1$ and $t + \tau_K + 1$ by

$$\begin{aligned} KN_c + \sum_{j=0}^{i-1} N_c N_3(j) \\ + \sum_{j=1}^{K-i} N_c \min(\max(1, \frac{N_3(i-1)}{2}) 2^{j-1}, L_1), \end{aligned} \quad (70)$$

where $N_3(j) = N_3(t + \tau_j + 1)$. To obtain the above expression, we have used the fact that when Case 1 happens, according to the update rule, at the $i + 1$ round $N_3(i)$ becomes half of the previous value for N_3 , but keeps doubling for each following round. Recalling that (68) holds after the i round, and $N_3(j) \leq L_1$, we can use (70) to show that for $K > i + 1$, w.p.1,

$$R_K \geq \frac{2 + \sum_{j=1}^{K-i-1} \min(2^{j-1}, L_1)}{K + iL_1 + 2 + \sum_{j=1}^{K-i-1} \min(2^{j-1}, L_1)}. \quad (71)$$

For a given fixed i , the above fraction approaches $\frac{L_1}{1+L_1}$ as K approaches ∞ . Therefore, for any given positive ϵ_2 , we can choose K sufficiently large such that for all i with $1 \leq i \leq i_{\max}$, the above fraction is larger than $\frac{L_1}{1+L_1} - \epsilon_2$. Applying a similar argument to the second case, we can find a sufficiently large K such that the fraction of time over which the near optimal solution is used is larger than $\frac{L_1}{1+L_1} - \epsilon_2$. In addition, we can select i_{\max} such that for a given positive ζ_1

$$\sum_{i=1}^{i_{\max}} (1 - \delta)^{i-1} \delta > 1 - \zeta_1.$$

Hence, if K is sufficiently large, with probability larger than $1 - \zeta_1$, we have

$$R_K \geq \left(\frac{L_1}{1+L_1} - \epsilon_2 \right). \quad (72)$$

This is an interesting observation. Since the choices for ϵ_2 and ζ_1 are arbitrary, this observation implies that in the limit of large backlog vectors, the policy keeps the network operating at near optimal points for at least $\frac{L_1}{1+L_1}$ fraction of time. Hence, in the limit, at most only the time for selecting new values for N_1 and observing their performance is *wasted*, which constitutes $\frac{1}{1+L_1}$ fraction of total time. Recall that near optimality is defined in (68), and $\phi(t + \tau_0 + 1, \tilde{N}_1) = \tilde{\phi}(\mathbf{X}_{t+\tau_0+1})$, we therefore, as a result of the preceding inequality, expect the limiting scaling factor of the capacity region to be a function of $\tilde{\phi}(\mathbf{X})$, and be proportional to $\frac{L_1}{1+L_1}$.

Note that to obtain the above results, in particular those mentioned in Case 1 and Case 2, we assume that the inequalities in (54) and (55) hold for all K scheduling rounds after time t . Therefore, the above discussion for R_K holds only for the limiting case of $\varrho_\varphi = 0$. In the following, we extend the

preceding discussions for a realistic situation where $\varrho_\varphi > 0$, and obtain a general lower bound for R_K .

Discussion assuming $\varrho_\varphi > 0$: We start by assuming that

$$\|\mathbf{X}_t\| > M''_K > M'_K, \quad (73)$$

for a sufficiently large M''_K such that for a given C , $C \gg 1$, Statement 1 holds for all times t_1 and t_2 greater than $t-1$ and less than $t+\tau_{(K+2)C}+1$, and Statement 2 holds for all i with $0 \leq i \leq (K+2)C$. We partition the time between $t+\tau_0+1$ and $t+\tau_{(K+2)C}+1$ into a set of *periods*, where each period consists of several scheduling rounds. For the simplicity of discussion, we assume that the first period always starts at $t+\tau_0+1$.

Corresponding to each period, e.g. the j th period, we define a *positive* r.v. $i_{\delta,j}$. This r.v. takes value i , $i > 0$, if the following conditions are met. First, in the $i+1$ th round of the j th period, for the first time in that period the optimal value for N_1 is selected. Second, the inequality (54) holds for φ^r at the $i+1$ th round as well as (55) for φ at the i th round, both in the j th period. Third, i equals $C-1$ if the last two conditions do not hold for any of the second to the $(C-2)$ th rounds in the j th period. Recall that the optimal N_1 is chosen independently in each round with probability δ . Thus, using Remark 2 with K replaced with $(K+2)C$, we see that $i_{\delta,j}$ becomes a truncated geometric r.v. with success probability

$$\delta' = (1 - \varrho_\varphi)^2 \delta, \quad (74)$$

and with the property that

$$P(i_{\delta,j} = C-1) = 1 - \sum_{i=1}^{C-2} \delta' (1 - \delta')^{i-1}. \quad (75)$$

Similarly, corresponding to the j th period, we define a *non-negative* r.v. denoted by $i_{\varphi,j}$ that is zero if $i_{\delta,j} = C-1$, and otherwise, is the number of consecutive rounds immediately following the $i_{\delta,j}$ th round in the j th period for all of which the inequalities in (54) and (55) hold. Similar to $i_{\delta,j}$, we limit $i_{\varphi,j}$ to be upper-bounded by $C-1$. We do so by letting $i_{\varphi,j} = C-1$ if for all $C-1$ rounds after the $i_{\delta,j}$ th round (54) and (55) hold. Using this definition of $i_{\varphi,j}$, and Remark 2 with K replaced with $(K+2)C$, it is easy to see that

$$P(i_{\varphi,j} = 0 | i_{\delta,j} = C-1) = 1, \quad (76)$$

$$P(i_{\varphi,j} = 0 | i_{\delta,j} \neq C-1) = \varrho_\varphi \quad (77)$$

and

$$\begin{aligned} P(i_{\varphi,j} = k | i_{\delta,j} \neq C-1) \\ = (1 - \varrho_\varphi)^{2k-1} (1 - (1 - \varrho_\varphi)^2), \quad 1 \leq k \leq C-2, \end{aligned} \quad (78)$$

and by the boundedness of $i_{\varphi,j}$,

$$\begin{aligned} P(i_{\varphi,j} = C-1 | i_{\delta,j} \neq C-1) \\ = 1 - \sum_{k=0}^{C-2} P(i_{\varphi,j} = k | i_{\delta,j} \neq C-1) \\ = (1 - \varrho_\varphi)^{2(C-1)-1}. \end{aligned} \quad (79)$$

To complete the characterization of periods, we define the last round in the j th period to be the one immediately following

the $i_{\delta,j} + i_{\varphi,j}$ th round in the j th period. This indicates that the j th period consists of $i_{\delta,j} + i_{\varphi,j} + 1$ rounds, and thus by the definition of $i_{\delta,j}$ and $i_{\varphi,j}$, its length is always less than $2C$.

Having introduced periods, we now define the sequence $\{p_j\}_{j=0}^\infty$, with $p_0 = 0$, as a subset of indices such that τ_{p_j} , $j \geq 1$, is the number of timeslots from time t to the last timeslot in the j th period. By definition, therefore, the j th period, $j \geq 1$, starts at $t + \tau_{p_{j-1}} + 1$ and ends at $t + \tau_{p_j} + 1$. Let i_K be number of periods that are completely contained in the K rounds after time t , i.e.,

$$i_K = \max\{j : p_j < K, j \geq 0\}. \quad (80)$$

By virtue of the definitions for a scheduling period, $i_{\delta,j}$, and $i_{\varphi,j}$, we can see that for all rounds after $i_{\delta,j}$ th and before the last round in the j th period, all conditions to apply the discussions in Case 1 and Case 2 are met. Hence, considering (68), for $1 \leq j \leq i_K + 1$ and $i_{\delta,j} < i \leq i_{\delta,j} + i_{\varphi,j}$ with $p_{j-1} + i - 1 < K$ we have that

$$|\varphi(t + \tau_{p_{j-1}+i-1} + 1) - \tilde{\phi}(t + \tau_0 + 1)| < \gamma. \quad (81)$$

Note that $t + \tau_{p_{j-1}+i-1} + 1$ is the start point of the i th round in the j th period, and we have set condition $p_{j-1} + i - 1 < K$ to consider only the first K rounds after time t .

We now focus on finding a lower bound for R_K . Towards that goal, we use r.v.'s $i_{\delta,j}$ and $i_{\varphi,j}$ to define a new sequence of N_3 denoted by N'_3 according to the following:

$$N'_3(k = p_{j-1} + i - 1) = N'_3(j, i) \\ = \begin{cases} L_1 & (1 \leq i \leq i_{\delta,j}) \vee (i = i_j + 1) \\ 1 & (i = i_{\delta,j} + 1) \wedge (i_{\varphi,j} = 1) \\ 2 & (i = i_{\delta,j} + 1) \wedge (i_{\varphi,j} > 1) \\ \min\left(\frac{2^i}{2^{i_{\delta,j}+2}}, L_1\right) & (i_{\delta,j} + 2 \leq i \leq i_j) \wedge (i_{\varphi,j} > 1) \end{cases}$$

where

$$i_j = i_{\delta,j} + i_{\varphi,j}.$$

Note that a round after time t can be specified uniquely either as the k th round after time t , or as the i th round in the j th period. We thus in the above have defined N'_3 as a function of the round number k after time t , and also as a function of the pair (j, i) . Similarly, N_3 can be considered as a function of either K or (j, i) . In addition, note that the above definition of N'_3 is mainly motivated by the method used to obtain (71).

To simplify the analysis, we slightly modify the definition of R_K such that

$$R_K = \frac{\sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} N_c N_3(j, i)}{\tau_K - \tau_0}.$$

Hence, R_K concerns only the rounds that are within the first i_K periods, and for which (81) holds. Considering the above definition, we can use a simple inspection to show that the above choices for N'_3 ensure that w.p.1

$$R_K \geq R'_K = \frac{\sum_{j=1}^{i_K} \lambda_r^C(j)}{\sum_{j=1}^{i_K+1} \lambda_t^C(j)}, \quad (82)$$

where

$$\begin{aligned}\lambda_t^C(j) &= \sum_{i=1}^{i_j+1} (1 + N_3'(j, i)) \\ &= i_j + 1 + (i_{\delta,j} + 1)L_1 \\ &+ \mathbf{1}_{(i_{\varphi,j}=1)} + (2 + \sum_{i=0}^{i_{\varphi,j}-2} \min(2^i, L_1)) \mathbf{1}_{(i_{\varphi,j}>1)},\end{aligned}\quad (83)$$

and

$$\begin{aligned}\lambda_r^C(j) &= \sum_{i=i_{\delta',j}+1}^{i_j} N_3'(j, i) \\ &= \mathbf{1}_{(i_{\varphi,j}=1)} + (2 + \sum_{i=0}^{i_{\varphi,j}-2} \min(2^i, L_1)) \mathbf{1}_{(i_{\varphi,j}>1)}.\end{aligned}\quad (84)$$

As expected, $\lambda_r^C(j)$ denotes the minimum contribution of the j th period to the ratio R_K . The term $\lambda_t^C(j)$ is the total length of the j th period that could potentially minimize the ratio R_K . In addition, note that the inequality (82) in general holds even when Remark 1 does not hold, and thus, when the distribution of $i_{\delta,j}$ and $i_{\varphi,j}$ is not given by (74)-(79). However, as stated in Remark 1, we consider the worst case which enables us to find a lower-bound for R_K that holds with high probability. We next show that the random variable R'_K is a function of i.i.d pairs, and in fact, is the average accumulated reward for a renewal process.

First, note that by definition $i_{\delta,j} > 0$, and hence, a scheduling period, which consists of $i_{\delta,j} + i_{\varphi,j} + 1$ rounds, at least contains of two rounds. This implies that the K rounds under consideration constitute at most $\lfloor \frac{K}{2} \rfloor$ complete periods. Consequently, R'_K is a function of at most $K_P = \lfloor \frac{K}{2} \rfloor + 1$ periods, and thus, is completely characterized by

$$\{(i_{\delta,j}, i_{\varphi,j}), 1 \leq j \leq K_P\}. \quad (85)$$

We know that by definition a period consists of at most $2C - 1$ rounds. Therefore, considering Remark 2 with K replaced with $K_P(2C - 1)$, we see that the above set is consisted of i.i.d pairs, with distribution given by (74)-(79), if Statement 2 holds for all i with $0 \leq i \leq K_P(2C - 1)$. Recall that we started by assuming $\|\mathbf{X}_t\| > M''_K$ such that Statement 2 holds for $0 \leq i \leq (K+2)C$. But this means that Statement 2 holds for all i with $0 \leq i \leq K_P(2C - 1)$ since $K_P(2C - 1) < (K+2)C$. Therefore, we have that the pairs in (85) are i.i.d.¹³

Next, observe that since the pair $(\lambda_t^C(j), \lambda_r^C(j))$ depends only on $(i_{\delta,j}, i_{\varphi,j})$, the sequence $\{(\lambda_t^C(j), \lambda_r^C(j)) : 1 \leq j \leq K_P\}$ also consists of i.i.d. pairs. This sequence is defined for $1 \leq j \leq K_P$, but can be defined for $j > K_P$ by defining the pair $(\lambda_t^C(j), \lambda_r^C(j))$, for $j > K_P$, as an i.i.d. version of $(\lambda_t^C(1), \lambda_r^C(1))$. The resulting expanded sequence

$$\{(\lambda_t^C(j), \lambda_r^C(j)) : j \geq 1\}$$

¹³Note that if $C = \infty$, $i_{\delta,j}$ or $i_{\varphi,j}$ may take any finite value. Hence, a proper definition of $i_{\delta,j}$ or $i_{\varphi,j}$ with distributions given by (74)-(79) requires Statement 2 hold for all $i \geq 0$, which cannot be true by assuming $\|\mathbf{X}_t\| > M''_K$, for any finite value of M''_K .

defines a reward renewal process. For this renewal process, $\lambda_t^C(j)$ is the length of the j th inter-renewal interval, $\lambda_r^C(j)$ is the accumulated reward collected at the end of j th renewal interval, and R'_K is the average accumulated reward prior to end of $i_K + 1$ th inter-renewal interval.

Consider the extended sequence, and let R'_k , for any $k > 0$, be defined similar to R'_K . Applying the strong law for the renewal process, and noting that $i_k \rightarrow \infty$, a.s., as $k \rightarrow \infty$, we obtain

$$R_\infty^C \triangleq \lim_{k \rightarrow \infty} R'_k = \frac{\mathbb{E}[\lambda_r^C(1)]}{\mathbb{E}[\lambda_t^C(1)]}, \quad \text{a.s.}$$

Hence, by the almost surely convergence, for any given $\epsilon_R > 0$ and $\varrho_R > 0$, there exists a sufficiently large $n_{\epsilon_R, \varrho_R}^C$ such that [29]

$$P\left(\sup_{k \geq n_{\epsilon_R, \varrho_R}^C} |R'_k - R_\infty^C| < \frac{\epsilon_R}{2}\right) > (1 - \varrho_R). \quad (86)$$

But since $\lim_{C \rightarrow \infty} R_\infty^C = R_\infty$, we can chose a sufficiently large C such that

$$|R_\infty^C - R_\infty| < \frac{\epsilon_R}{2}.$$

Considering (86) for this value of C , we have that

$$P\left(\sup_{k \geq n_{\epsilon_R, \varrho_R}^C} |R'_k - R_\infty| < \epsilon_R\right) > (1 - \varrho_R). \quad (87)$$

The above inequality and (82) imply that there exists a sufficiently large $K_{\epsilon_R, \varrho_R}$ such that for $K > K_{\epsilon_R, \varrho_R}$ and $\|\mathbf{X}_t\| > M''_K$

$$P(R_K > R_\infty - \epsilon_R | \mathcal{H}_{t+\tau_0+1}) > (1 - \varrho_R). \quad (88)$$

Here, we have stated the probability conditioned on $\mathcal{H}_{t+\tau_0+1}$ since all of the previous discussions are valid regardless of $\mathcal{H}_{t+\tau_0+1}$. The above inequality states that with probability close to one, R_K is close to R_∞ in the sense that $R_K > R_\infty - \epsilon_R$. This is a generalized version of the result obtained in (72), as desired.

We are finally in a position to derive a lower bound for the LHS of the inequality in the lemma, denoted by Σ . First, note that

$$\begin{aligned}\Sigma &= \mathbb{E}\left[\sum_{i=\tau_0+1}^{\tau_K} \mathbf{X}_{t+i} \mathbf{D}_{t+i} | \mathcal{H}_t\right] \\ &\geq \mathbb{E}\left[\sum_{k=1}^{p_{i_K}} \sum_{i=1}^{\tau_K - \tau_{k-1} - N_c} \mathbf{X}_{t+\tau_{k-1} + N_c + i} \mathbf{D}_{t+\tau_{k-1} + N_c + i} | \mathcal{H}_t\right],\end{aligned}\quad (89)$$

where we have simply used the fact that the product $\mathbf{X}_{t+i} \mathbf{D}_{t+i}$ is positive, and neglected the contributions due to the test intervals, and also the ones due to the rounds of the last partially covered period.

To simplify the notation, let $t_{j,i,l}$ denote the start of l th timeslot of the i th round in the j th period, i.e.,

$$t_{j,i,l} = t + \tau_{p_{j-1}+i-1} + l.$$

In addition, let $\delta_{j,i}$ denote the length of the i th round in the j th period, i.e.,

$$\delta_{j,i} = \tau_{p_{j-1}+i} - \tau_{p_{j-1}+i-1} = N_c(1 + N_3(j, i)),$$

where

$$N_3(j, i) = N_3(t + \tau_{p_{j-1}+i-1} + 1).$$

Considering the partition generated by the scheduling periods, and the above definitions, we can use (89) to show that

$$\begin{aligned} \Sigma &\geq \mathbb{E} \left[\sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} \sum_{l=1}^{\delta_{j,i}-N_c} \mathbf{X}_{t_{j,i,N_c+l}} \mathbf{D}_{t_{j,i,N_c+l}} | \mathcal{H}_t \right] \\ &\geq \mathbb{E} \left[\sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} \sum_{l=1}^{\delta_{j,i}-N_c} \|\mathbf{X}_t\| \right. \\ &\quad \left. \frac{\|\mathbf{X}_{t_{j,i,N_c+l}}\|}{\|\mathbf{X}_t\|} (\tilde{\phi}(t + \tau_0 + 1) - \gamma) | \mathcal{H}_t \right], \end{aligned} \quad (90)$$

where the last inequality follows from (81). Using (57) and assuming

$$\frac{\|\mathbf{X}_{t_{j,i,N_c+l}}\|}{\|\mathbf{X}_t\|} > (1 - \epsilon_3),$$

we obtain

$$\begin{aligned} \Sigma &\geq \mathbb{E} \left[\sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} \sum_{l=1}^{\delta_{j,i}-N_c} \|\mathbf{X}_t\| (1 - \epsilon_3) (\tilde{\phi}(t) - \epsilon_1 - \gamma) | \mathcal{H}_t \right] \\ &= \mathbb{E} \left[\|\mathbf{X}_t\| (\tilde{\phi}(t) - \epsilon_1 - \gamma) \right. \\ &\quad \left. (1 - \epsilon_3) \sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} N_c N_3(j, i) | \mathcal{H}_t \right]. \end{aligned} \quad (91)$$

But, by using (88) and adopting a method similar to the one in Lemma 1, for $K > K_{\epsilon_R, \varrho_R}$ and $\|\mathbf{X}_t\| > M''_K$ we can show that

$$\begin{aligned} \mathbb{E} \left[\sum_{j=1}^{i_K} \sum_{i=i_{\delta,j}+1}^{i_j} N_c N_3(j, i) | \mathcal{H}_{t+\tau_0+1} \right] \\ \geq (1 - \epsilon_4) (R_\infty - \epsilon_R) \mathbb{E} \left[(\tau_K - \tau_0) | \mathcal{H}_{t+\tau_0+1} \right], \end{aligned} \quad (92)$$

where $\epsilon_4 \rightarrow 0$, as $K \rightarrow \infty$.

Using (91) and (92), we obtain

$$\begin{aligned} \Sigma &\geq \mathbb{E} \left[(\tau_K - \tau_0) \|\mathbf{X}_t\| (\tilde{\phi}(t) - \epsilon_1 - \gamma) \right. \\ &\quad \left. (R_\infty - \epsilon_R) (1 - \epsilon_3) (1 - \epsilon_4) | \mathcal{H}_t \right]. \end{aligned}$$

If we assume

$$\frac{\tau_K - \tau_0}{\tau_K + 1} > 1 - \epsilon_5,$$

then using the above inequality and the definition of γ , given in (69), we have that

$$\Sigma \geq \mathbb{E} \left[(\tau_K + 1) \|\mathbf{X}_t\| (R_\infty (\tilde{\phi}(t) - \alpha - 3\theta_\varphi) - \epsilon) | \mathcal{H}_t \right],$$

where $\epsilon > 0$, and can be made arbitrarily small by choosing sufficiently small values for $\epsilon_1, \epsilon_3, \epsilon_4, \epsilon_5$, and ϵ_R . Note that since $\|\mathbf{X}_t - \mathbf{X}_{t+i}\| < C'_K$, $0 \leq i \leq \tau_K$, as discussed in the beginning of the proof of the lemma, ϵ_3 can be assumed arbitrarily small if $\|\mathbf{X}_t\|$ is sufficiently large. Moreover, since $\tau_0 \leq (1 + L_1)N_c$ and $\tau_K + 1 \geq 2KN_c$, ϵ_5 can be made arbitrarily small by assuming a sufficiently large K . Thus, by

considering the discussions for ϵ_R and ϵ_4 , we see that we can make ϵ_R, ϵ_4 , and ϵ_5 all sufficiently small by choosing $K > K_\epsilon$, for sufficiently large K_ϵ . Having selected K , we can find a lower bound $M_{\epsilon, K} > M''_K > M'_K$ for $\|\mathbf{X}_t\|$ such that ϵ_1 and ϵ_3 are also sufficiently small. Hence, ϵ can be arbitrarily small, completing the proof of the lemma. \blacksquare

Lemma 5: Let $0 \leq \delta < 1$. We have

$$\prod_{i=1}^{\infty} (1 - \delta^i) > \exp \left(- \left(\frac{\delta}{(1 - \delta)} + \frac{\delta^2}{(1 - \delta)^2(1 - \delta^2)} \right) \right) > 0.$$

Proof: First note that by Taylor's theorem, we have

$$\ln(1 - \delta) \geq -\left(\delta + \frac{\delta^2}{2(1 - \delta)^2}\right).$$

Taking \ln and then \exp of the product term in the lemma, and using the above inequality, we can easily show that

$$\begin{aligned} \prod_{i=1}^{\infty} (1 - \delta^i) &\geq \exp \left(- \sum_{i=1}^{\infty} \delta^i + \sum_{i=1}^{\infty} \frac{\delta^{2i}}{(1 - \delta)^2} \right) \\ &= \exp \left(- \left(\frac{\delta}{(1 - \delta)} + \frac{\delta^2}{(1 - \delta)^2(1 - \delta^2)} \right) \right) > 0, \end{aligned}$$

proving the lemma. \blacksquare

REFERENCES

- [1] M. Lotfinezhad, B. Liang, and E. Sousa, "Dynamic control of tunable sub-optimal algorithms for scheduling of time-varying wireless networks," in *Sixteenth IEEE International Workshop on Quality of Service (IWQoS)'08*, June 2008.
- [2] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [3] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 89–103, Jan. 2005.
- [4] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 411–424, Apr. 2005.
- [5] X. Lin and N. Shroff, "The impact of imperfect scheduling on cross-layer congestion control in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 302–315, Apr. 2006.
- [6] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *Proc. of the 12th annual international conference on Mobile computing and networking (MobiCom'06)*, 2006, pp. 227–238.
- [7] X. Wu and R. Srikant, "Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing," in *44th IEEE conference on decision and control, and European control conference CDC-ECC*, Dec. 2005.
- [8] M. Lotfinezhad, B. Liang, and E. Sousa, "On stability region and delay performance of linear-memory randomized scheduling for time-varying networks," *To appear in IEEE/ACM Transactions on Networking*, 2009.
- [9] G. Sharma, N. B. Shroff, and R. R. Mazumdar, "Joint congestion control and distributed scheduling for throughput guarantees in wireless networks," in *Proc. IEEE INFOCOM'07*, May 2007.
- [10] B. H. Fralix, "Foster-type criteria for Markov chains on general spaces," *Journal of Applied Probability*, vol. 43, no. 4, pp. 1194–1200, 2006.
- [11] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," *Translations of AMS, A Volume in Memory of F. Karpelevich*, vol. 207, 2002.
- [12] E. Yeh and R. Berry, "Throughput optimal control of cooperative relay networks," *IEEE Trans. Inf. Theory*, vol. 53, no. 10, pp. 3827–3833, Oct. 2007.
- [13] M. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 396–409, Apr. 2008.

- [14] M. Neely, "Optimal energy and delay tradeoffs for multiuser wireless downlinks," *IEEE Trans. Inf. Theory*, vol. 53, no. 9, pp. 3095–3113, Sept. 2007.
- [15] W. Chen, M. Neely, and U. Mitra, "Energy-efficient transmissions with individual packet delay constraints," *IEEE Trans. Inf. Theory*, vol. 54, no. 5, pp. 2090–2109, May 2008.
- [16] A. Eryilmaz and R. Srikant, "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," *IEEE/ACM Trans. Netw.*, vol. 15, no. 6, pp. 1333–1344, Dec. 2007.
- [17] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, "Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks," in *Proc. IEEE INFOCOM'06*, Apr. 2006.
- [18] M. J. Neely, "Optimal pricing in a free market wireless network," in *Proc. IEEE INFOCOM'07*, May 2007.
- [19] P. Marbach, "Distributed scheduling and active queue management in wireless networks," in *Proc. IEEE INFOCOM'07*, May 2007.
- [20] L. Jiang and J. Walrand, "A distributed algorithm for maximal throughput and optimal fairness in wireless networks with a general interference model," EECS Department, University of California, Berkeley, Tech. Rep., Apr 2008. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-38.html>
- [21] M. Goyal, A. Kumar, and V. Sharma, "Optimal cross-layer scheduling of transmissions over a fading multiaccess channel," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3518–3537, Aug. 2008.
- [22] L. Tassiulas, "Linear complexity algorithms for maximum throughput in radio networks and input queued switches," in *Proc. IEEE INFOCOM'98*, vol. 2, Mar./Apr. 1998, pp. 533–539.
- [23] E. Modiano, D. Shah, and G. Zussman, "Maximizing throughput in wireless networks via gossiping," *SIGMETRICS Perform. Eval. Rev.*, vol. 34, no. 1, pp. 27–38, 2006.
- [24] P. Chaporkar and S. Sarkar, "Stable scheduling policies for maximizing throughput in generalized constrained queueing systems," *IEEE Trans. Autom. Control*, vol. 53, no. 8, pp. 1913–1931, Sep. 2008.
- [25] X. Wu, R. Srikant, and J. R. Perkins, "Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks," *IEEE Trans. Mobile Comput.*, vol. 6, no. 6, pp. 595–605, Jun. 2007.
- [26] M. J. Neely, "Delay analysis for maximal scheduling in wireless networks with bursty traffic," in *Proc. IEEE INFOCOM'08*, Apr. 2008.
- [27] P. Chaporkar, K. Kar, X. Luo, and S. Sarkar, "Throughput and fairness guarantees through maximal scheduling in wireless networks," *IEEE Trans. Inf. Theory*, vol. 54, no. 2, pp. 572–594, Feb. 2008.
- [28] C. Joo, X. Lin, and N. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks," in *Proc. IEEE INFOCOM'08*, Apr. 2008.
- [29] A. A. Borovkov, *Probability Theory*. Gordon and Breach, 1998.
- [30] L. Tassiulas, "Scheduling and performance limits of networks with constantly changing topology," *IEEE Trans. Inf. Theory*, vol. 43, no. 3, pp. 1067–1073, May 1997.
- [31] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [32] R. Buche and H. J. Kushner, "Control of mobile communication systems with time-varying channels via stability methods," *IEEE Trans. Autom. Control*, vol. 49, no. 11, pp. 1954–1962, Nov. 2004.
- [33] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [34] J. S. Rosenthal, *A First Look at Rigorous Probability Theory*. World Scientific Publishing Co., 2003.